

**IN THE UNITED STATES DISTRICT COURT
FOR DISTRICT OF DELAWARE**

ORCA SECURITY LTD.,)	
)	
Plaintiff,)	C.A. No. 23-0758-JLH
)	
v.)	
)	
WIZ, INC.)	
)	
Defendant.)	

**DEFENDANT WIZ INC.’S ANSWER TO SECOND AMENDED COMPLAINT AND
COUNTERCLAIMS**

Defendant and Counterclaim-Plaintiff Wiz, Inc. (“Wiz”) hereby responds to the Second Amended Complaint (“Complaint”) for patent infringement (D.I. 15) of Plaintiff and Counterclaim-Defendant Orca Security Ltd. (“Orca”) as follows. To the extent not specifically admitted in the following paragraphs, the allegations in Orca’s Second Amended Complaint are denied.

INTRODUCTION AND SUMMARY OF THE ACTION¹

1. Wiz denies the allegations in paragraph 1 of the Complaint.
2. Wiz denies the allegations in paragraph 2 of the Complaint.
3. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 3 of the Complaint and on that basis, denies them.

¹ Wiz has incorporated the headings that appear in the Second Amended Complaint. Wiz does not necessarily agree with the characterization of such headings and does not waive any right to object to those characterizations. Accordingly, to the extent that a particular heading can be construed as an allegation, Wiz specifically denies any such allegations.

4. Plaintiff's allegations in paragraph 4 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of Plaintiff's patents.

5. Plaintiff's allegations in paragraph 5 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of Plaintiff's patents.

6. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 6 of the Complaint and on that basis, denies them.

7. Plaintiff's allegations in paragraph 7 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of Plaintiff's patents.

8. Plaintiff's allegations in paragraph 8 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of Plaintiff's patents.

9. Wiz specifically denies that it infringes any valid claim of Plaintiff's patents. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 9 of the Complaint and on that basis, denies them.

10. Wiz admits that the faces of what appear to be U.S. Patent Nos. 11,663,031 (the "'031 patent"), 11,663,032 (the "'032 patent"), 11,693,685 (the "'685 patent"), 11,726,809 (the

“‘809 patent”), 11,740,926 (the “‘926 patent”), and 11,775,326 (the “‘326 patent”) list “Avi Shua” as “Inventor.” Wiz denies that these patents issued on August 22, 2022, and specifically denies that Wiz infringes any valid claim of Plaintiff’s patents.² Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 10 of the Complaint and on that basis, denies them.

11. Wiz denies the allegations in paragraph 11 of the Complaint.

12. Wiz denies the allegations in paragraph 12 of the Complaint.

WIZ AND ITS WIDESPREAD COPYING OF ORCA

13. Wiz admits that it was founded in January 2020 by Messrs. Rappaport, Luttwak, Costica, and Reznik. Wiz further admits that Messrs. Rappaport, Luttwak, Costica, and Reznik joined Microsoft after Microsoft acquired their prior cloud security company, Adallom, Inc. (“Adallom”) and that they served as leaders of Microsoft’s Cloud Security Group. Wiz further admits that, after years in the security industry at Adallom and Microsoft, the founders of Wiz identified prior cloud security solutions as complex, fragmented, and generating too many alerts for security teams. Wiz further admits that on or around December 9, 2020, Wiz emerged from stealth with a cloud security solution that took a new approach and with a new architecture allowing for seamless scanning of the entire cloud environment across compute types and cloud services for vulnerabilities, configuration, network, and identity issues without agents or sidecars. The remaining allegations of paragraph 13 appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the

² Orca is likely referring to U.S. Patent No. 11,431,735, which was the first patent to issue in the family of asserted patents and issued on August 22, 2022. In response to a petition for inter partes review by Wiz, Orca statutorily disclaimed all challenged claims in the ‘735 patent, including all independent claims. See IPR2024-00220, Doc. No. 6.

claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of Plaintiff's patents.

14. Wiz admits that Messrs. Rappaport, Luttwak, Costica, and Reznik left Microsoft at various dates and thereafter founded Wiz in January 2020. Wiz denies the remaining allegations in paragraph 14 of the Complaint, and specifically denies that it infringes any valid claim of Plaintiff's patents or engaged in any actionable copying of Plaintiff.

15. Wiz admits that by in or around December 2020, it had a cloud security solution that allowed for seamless scanning of the entire cloud environment across compute types and cloud services for vulnerabilities, configuration, network, and identity issues without agents or sidecars and that the solution delivered 360 degrees of visibility for cloud security teams by highlighting the critical risks in cloud environments across all risk pillars, and with the goal to empower security teams to know their clouds better than the developer teams. Wiz further admits that Fortune 100 companies were among its early customers. Wiz further admits that it issued the press release available at <https://www.wiz.io/blog/100m-arr-in-18-months-wiz-becomes-the-fastest-growing-software-company-ever> on or around August 10, 2022. The press release speaks for itself. Wiz further admits that by in or around February 2023, it had raised \$300 million at a \$10 billion valuation. Wiz denies the remaining allegations in paragraph 15 of the Complaint and specifically denies that it infringes any valid claim of Plaintiff's patents or engaged in any actionable copying of Plaintiff.

16. Wiz denies that it has engaged in any copying of "Orca's technology." Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 16 of the Complaint and on that basis, denies them.

17. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 17 of the Complaint and on that basis, denies them.

18. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 18 of the Complaint and on that basis, denies them.

19. The allegations of paragraph 19 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies that it infringes any valid claim of Plaintiff's patents or engaged in any actionable copying of Plaintiff. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 19 of the Complaint and on that basis, denies them.

20. The allegations of paragraph 20 include purported citation to a document from Wiz's website; that document speaks for itself. Wiz denies that it infringes any valid claim of Plaintiff's patents or engaged in any actionable copying of Plaintiff. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 20 of the Complaint and on that basis, denies them.

21. The allegations of paragraph 21 include purported citation to Wiz's website; that website speaks for itself. Wiz denies that it infringes any valid claim of Plaintiff's patents or engaged in any actionable copying of Plaintiff. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 21 of the Complaint and on that basis, denies them.

22. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegation that "Wiz showed up with its own coffee machine," and, on that basis, denies it. Wiz denies the remaining allegations of paragraph 22 of the Complaint, and

specifically denies that it infringes any valid claim of Plaintiff's patents or engaged in any actionable copying of Plaintiff.

23. Wiz admits that it was issued U.S. Patent No. 11,374,982; that document speaks for itself. Wiz denies copying Orca's patents, its prosecution strategy, or its prosecuting attorney. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 23 of the Complaint and on that basis, denies them.

24. Wiz denies the allegations in paragraph 24 of the Complaint.

25. Wiz denies hiring Orca's outside corporate counsel to assist Wiz in an attempt to copy Orca. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the remaining allegations in paragraph 25 of the Complaint and on that basis, denies them.

26. Wiz denies the allegations in paragraph 26 of the Complaint.

27. Wiz denies the allegations in paragraph 27 of the Complaint.

28. Wiz denies the allegations in paragraph 28 of the Complaint.

29. Wiz denies the allegations in paragraph 29 of the Complaint.

THE PARTIES

30. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 30 of the Complaint and on that basis, denies them.

31. Wiz admits that Wiz, Inc. is a Delaware company with a principal place of business at One Manhattan West, 57th Floor, New York, New York.

JURISDICTION AND VENUE

32. Wiz admits that this action invokes the United States patent laws, and that this Court has subject matter jurisdiction over patent law claims pursuant to 28 U.S.C. §§ 1331 and 1338(a).

33. Wiz does not contest that this Court has personal jurisdiction solely for the purposes of this particular action. Wiz specifically denies that it has committed any acts of infringement within this district, or any other district. Otherwise denied.

34. Wiz does not contest that this Court has personal jurisdiction solely for the purposes of this particular action. Wiz specifically denies that it has committed any acts of infringement within this district, or any other district. Otherwise denied.

35. Wiz admits that venue is proper in this judicial district for the purposes of this particular action. Wiz specifically denies that it has committed any acts of infringement within this district, or any other district. Otherwise denied.

COUNT I
(INFRINGEMENT OF THE '031 PATENT)

36. Wiz realleges and incorporates by reference its responses to paragraphs 1-35.

37. Wiz admits that what purports to be a copy of U.S. Patent No. 11,663,031 is attached as Exhibit 1 to the Complaint. Wiz further admits that the face of what appears to be the '031 patent indicates that its title is “Techniques for Securing Virtual Cloud Assets at Rest Against Cyber Threats” and that the date of the patent is May 30, 2023. Wiz denies that the '031 patent was duly and legally issued.

38. Wiz denies that Orca has any right to recover damages for infringement of the '031 patent. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 38 of the Complaint and on that basis, denies them.

39. Wiz denies that the '031 patent is valid and enforceable.

40. Orca's allegations in paragraph 40 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '031 patent.

41. Wiz denies the allegations in paragraph 41 of the Complaint.

42. Wiz admits that paragraph 42 of the Complaint reproduces the language of claim 9 of what appears to be the '031 patent.

43. Wiz denies the allegations in paragraph 43 of the Complaint.

44. Wiz admits the claim language quoted in paragraph 44 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 44 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

45. Wiz admits the claim language quoted in paragraph 45 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 45 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

46. Wiz admits the claim language quoted in paragraph 46 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 46 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

47. Wiz admits the claim language quoted in paragraph 47 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in

paragraph 47 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

48. Wiz admits the claim language quoted in paragraph 48 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 48 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

49. Wiz admits the claim language quoted in paragraph 49 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 49 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

50. Wiz admits the claim language quoted in paragraph 50 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 50 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

51. Wiz admits the claim language in paragraph 51 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 51 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

52. Orca's allegations in paragraph 52 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '031 patent.

53. Wiz admits the claim language in paragraph 53 of the Complaint appears in claim 9 of what appears to be the '031 patent. Wiz denies the remaining allegations in paragraph 53 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

54. Wiz denies the allegations in paragraph 54 of the Complaint.

55. Wiz denies the allegations in paragraph 55 of the Complaint.

56. Wiz denies the allegations in paragraph 56 of the Complaint.

57. Wiz denies the allegations in paragraph 57 of the Complaint.

58. Wiz denies the allegations in paragraph 58 of the Complaint.

59. The allegations in paragraph 59 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 59 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

60. The allegations in paragraph 60 include purported citation to a video posted by Wiz; that video speaks for itself. Wiz denies the remaining allegations in paragraph 60 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

61. The allegations in paragraph 61 include purported citation to documents from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 61 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

62. Wiz denies the allegations in paragraph 62 of the Complaint.

63. Wiz denies the allegations in paragraph 63 of the Complaint.

64. Wiz denies the allegations in paragraph 64 of the Complaint.

65. The allegations in paragraph 65 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 65 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

66. Wiz denies the allegations in paragraph 66 of the Complaint.

67. The allegations in paragraph 67 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 67 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

68. The allegations in paragraph 68 include purported citation to a video posted by Wiz; that video speaks for itself. Wiz denies the remaining allegations in paragraph 68 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

69. The allegations in paragraph 69 include purported citation to documents from Wiz's website; that website speaks for itself. Wiz denies the allegations in paragraph 69 of the Complaint, and specifically denies that it infringes any valid claim of the '031 patent.

70. Wiz denies the allegations in paragraph 70 of the Complaint.

71. Wiz denies the allegations in paragraph 71 of the Complaint.

72. Wiz denies the allegations in paragraph 72 of the Complaint.

COUNT II
(INFRINGEMENT OF THE '032 PATENT)

73. Wiz realleges and incorporates by reference its responses to paragraphs 1-35.

74. Wiz admits that what purports to be a copy of U.S. Patent No. 11,663,032 is attached as Exhibit 2 to the Complaint. Wiz further admits that the face of what appears to be the '032 patent indicates that its title is "Techniques for Securing Virtual Machines By Application Use Analysis" and that the date of the patent is May 30, 2023. Wiz denies that the '032 patent was duly and legally issued.

75. Wiz denies that Orca has any right to recover damages for infringement of the '032 patent. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 75 of the Complaint and on that basis, denies them.

76. Wiz denies that the '032 patent is valid and enforceable.

77. Orca's allegations in paragraph 77 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '032 patent.

78. Orca's allegations in paragraph 78 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '032 patent.

79. Wiz denies the allegations in paragraph 79 of the Complaint.

80. Wiz admits that paragraph 80 of the Complaint reproduces the language of claim 1 of what appears to be the '032 patent.

81. Wiz denies the allegations in paragraph 81 of the Complaint.

82. Wiz admits the claim language quoted in paragraph 82 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in paragraph 82 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

83. Wiz admits the claim language quoted in paragraph 83 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in paragraph 83 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

84. Wiz admits the claim language quoted in paragraph 84 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in

paragraph 84 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

85. Wiz admits the claim language quoted in paragraph 85 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in paragraph 85 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

86. Wiz admits the claim language quoted in paragraph 86 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in paragraph 86 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

87. Wiz admits the claim language quoted in paragraph 87 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in paragraph 87 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

88. Wiz admits the claim language quoted in paragraph 88 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in paragraph 88 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

89. Wiz admits the claim language quoted in paragraph 89 of the Complaint appears in claim 1 of what appears to be the '032 patent. Wiz denies the remaining allegations in paragraph 89 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

90. Wiz denies the allegations in paragraph 90 of the Complaint.

91. Wiz denies the allegations in paragraph 91 of the Complaint.

92. Wiz denies the allegations in paragraph 92 of the Complaint.

93. Wiz denies the allegations in paragraph 93 of the Complaint.

94. Wiz denies the allegations in paragraph 94 of the Complaint.

95. The allegations in paragraph 95 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 95 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

96. The allegations in paragraph 96 include purported citation to a video posted by Wiz; that video speaks for itself. Wiz denies the remaining allegations in paragraph 96 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

97. The allegations in paragraph 97 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 97 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

98. Wiz denies the allegations in paragraph 98 of the Complaint.

99. Wiz denies the allegations in paragraph 99 of the Complaint.

100. The allegations in paragraph 100 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 100 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

101. The allegations in paragraph 101 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 101 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

102. Wiz denies the allegations in paragraph 102 of the Complaint.

103. The allegations in paragraph 103 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 103 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

104. The allegations in paragraph 104 include purported citation to a video posted by Wiz; that video speaks for itself. Wiz denies the remaining allegations in paragraph 104 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

105. The allegations in paragraph 96 include purported citation to a video posted by Wiz; that video speaks for itself. Wiz denies the remaining allegations in paragraph 96 of the Complaint, and specifically denies that it infringes any valid claim of the '032 patent.

106. Wiz denies the allegations in paragraph 106 of the Complaint.

107. Wiz denies the allegations in paragraph 107 of the Complaint.

108. Wiz denies the allegations in paragraph 108 of the Complaint.

COUNT III
(INFRINGEMENT OF THE '685 PATENT)

109. Wiz realleges and incorporates by reference its responses to paragraphs 1-35.

110. Wiz admits that what purports to be a copy of U.S. Patent No. 11,693,685 is attached as Exhibit 7 to the Complaint. Wiz further admits that the face of what appears to be the '685 patent indicates that its title is "Virtual Machine Vulnerabilities and Sensitive Data Analysis and Detection" and that the date of the patent is July 4, 2023. Wiz denies that the '685 patent was duly and legally issued.

111. Wiz denies that Orca has any right to recover damages for infringement of the '685 patent. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 111 of the Complaint and on that basis, denies them.

112. Wiz denies that the '685 patent is valid and enforceable.

113. Orca's allegations in paragraph 113 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '685 patent.

114. Wiz denies the allegations in paragraph 114 of the Complaint.

115. Wiz admits that paragraph 115 of the Complaint reproduces the language of claim 1 of what appears to be the '685 patent.

116. Wiz denies the allegations in paragraph 116 of the Complaint.

117. Wiz admits the claim language quoted in paragraph 117 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 117 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

118. Wiz admits the claim language quoted in paragraph 118 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 118 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

119. Wiz admits the claim language quoted in paragraph 119 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 119 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

120. Wiz admits the claim language quoted in paragraph 120 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in

paragraph 120 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

121. Wiz admits the claim language quoted in paragraph 121 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 121 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

122. Wiz admits the claim language quoted in paragraph 122 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 122 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

123. Wiz admits the claim language quoted in paragraph 123 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 123 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

124. Wiz admits the claim language quoted in paragraph 124 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 124 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

125. Wiz admits the claim language quoted in paragraph 125 of the Complaint appears in claim 1 of what appears to be the '685 patent. Wiz denies the remaining allegations in paragraph 125 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

126. Wiz denies the allegations in paragraph 126 of the Complaint.

127. Wiz denies the allegations in paragraph 127 of the Complaint.

128. Wiz denies the allegations in paragraph 128 of the Complaint.

129. Wiz denies the allegations in paragraph 129 of the Complaint.

130. Wiz denies the allegations in paragraph 130 of the Complaint.

131. The allegations in paragraph 131 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 131 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

132. The allegations in paragraph 132 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 132 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

133. The allegations in paragraph 133 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 133 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

134. Wiz denies the allegations in paragraph 134 of the Complaint.

135. Wiz denies the allegations in paragraph 135 of the Complaint.

136. Wiz denies the allegations in paragraph 136 of the Complaint.

137. The allegations in paragraph 137 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 137 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

138. Wiz denies the allegations in paragraph 138 of the Complaint.

139. The allegations in paragraph 139 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 139 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

140. The allegations of paragraph 140 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 140 of the Complaint, and specifically denies that it infringes any valid claim of the '685 patent.

141. Wiz denies the allegations in paragraph 141 of the Complaint.

142. Wiz denies the allegations in paragraph 142 of the Complaint.

143. Wiz denies the allegations in paragraph 143 of the Complaint.

COUNT IV
(INFRINGEMENT OF THE '809 PATENT)

144. Wiz realleges and incorporates by reference its responses to paragraphs 1-35.

145. Wiz admits that what purports to be a copy of U.S. Patent No. 11,726,809 is attached as Exhibit 8 to the Complaint. Wiz further admits that the face of what appears to be the '809 patent indicates that its title is "Techniques for Securing Virtual Machines by Application Existence Analysis" and that the date of the patent is August 15, 2023. Wiz denies that the '809 patent was duly and legally issued.

146. Wiz denies that Orca has any right to recover damages for infringement of the '809 patent. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 146 of the Complaint and on that basis, denies them.

147. Wiz denies that the '809 patent is valid and enforceable.

148. Orca's allegations in paragraph 146 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '809 patent.

149. Wiz denies the allegations in paragraph 149 of the Complaint.

150. Wiz admits that paragraph 150 of the Complaint reproduces the language of claim 1 of what appears to be the '809 patent.

151. Wiz denies the allegations in paragraph 151 of the Complaint.

152. Wiz admits the claim language quoted in paragraph 152 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 152 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

153. Wiz admits the claim language quoted in paragraph 153 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 153 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

154. Wiz admits the claim language quoted in paragraph 154 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 154 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

155. Wiz admits the claim language quoted in paragraph 155 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 155 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

156. Wiz admits the claim language quoted in paragraph 156 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 156 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

157. Wiz admits the claim language quoted in paragraph 157 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 157 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

158. Wiz admits the claim language quoted in paragraph 158 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 158 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

159. Wiz admits the claim language quoted in paragraph 159 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 159 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

160. Wiz admits the claim language quoted in paragraph 160 of the Complaint appears in claim 1 of what appears to be the '809 patent. Wiz denies the remaining allegations in paragraph 160 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

161. Wiz denies the allegations in paragraph 161 of the Complaint.

162. Wiz denies the allegations in paragraph 162 of the Complaint.

163. Wiz denies the allegations in paragraph 163 of the Complaint.

164. Wiz denies the allegations in paragraph 164 of the Complaint.

165. Wiz denies the allegations in paragraph 165 of the Complaint.

166. The allegations in paragraph 166 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 166 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

167. The allegations in paragraph 167 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 167 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

168. The allegations in paragraph 168 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 168 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

169. Wiz denies the allegations in paragraph 169 of the Complaint.

170. Wiz denies the allegations in paragraph 170 of the Complaint.

171. Wiz denies the allegations in paragraph 171 of the Complaint.

172. The allegations in paragraph 172 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 172 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

173. Wiz denies the allegations in paragraph 173 of the Complaint.

174. The allegations in paragraph 174 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 174 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

175. The allegations in paragraph 175 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 175 of the Complaint, and specifically denies that it infringes any valid claim of the '809 patent.

176. Wiz denies the allegations in paragraph 176 of the Complaint.

177. Wiz denies the allegations in paragraph 177 of the Complaint.

178. Wiz denies the allegations in paragraph 178 of the Complaint.

COUNT V
(INFRINGEMENT OF THE '926 PATENT)

179. Wiz realleges and incorporates by reference its responses to paragraphs 1-35.

180. Wiz admits that what purports to be a copy of U.S. Patent No. 11,740,926 is attached as Exhibit 9 to the Complaint. Wiz further admits that the face of what appears to be the '926 patent indicates that its title is "Techniques for Securing Virtual Machines by Analyzing Data for Cyber Threats" and that the date of the patent is August 29, 2023. Wiz denies that the '926 patent was duly and legally issued.

181. Wiz denies that Orca has any right to recover damages for infringement of the '926 patent. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 181 of the Complaint and on that basis, denies them.

182. Wiz denies that the '809 patent is valid and enforceable.

183. Orca's allegations in paragraph 183 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '926 patent.

184. Wiz denies the allegations in paragraph 184 of the Complaint.

185. Wiz admits that paragraph 185 of the Complaint reproduces the language of claim 1 of what appears to be the '926 patent.

186. Wiz denies the allegations in paragraph 186 of the Complaint.

187. Wiz admits the claim language quoted in paragraph 187 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in

paragraph 187 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

188. Wiz admits the claim language quoted in paragraph 188 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in paragraph 188 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

189. Wiz admits the claim language quoted in paragraph 189 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in paragraph 189 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

190. Wiz admits the claim language quoted in paragraph 190 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in paragraph 190 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

191. Wiz admits the claim language quoted in paragraph 191 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in paragraph 191 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

192. Wiz admits the claim language quoted in paragraph 192 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in paragraph 192 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

193. Wiz admits the claim language quoted in paragraph 193 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in paragraph 193 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

194. Wiz admits the claim language quoted in paragraph 194 of the Complaint appears in claim 1 of what appears to be the '926 patent. Wiz denies the remaining allegations in paragraph 194 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

195. Wiz denies the allegations in paragraph 195 of the Complaint.

196. Wiz denies the allegations in paragraph 196 of the Complaint.

197. Wiz denies the allegations in paragraph 197 of the Complaint.

198. Wiz denies the allegations in paragraph 198 of the Complaint.

199. Wiz denies the allegations in paragraph 199 of the Complaint.

200. The allegations in paragraph 200 include purported citation to a video posted by Wiz; that video speaks for itself. Wiz denies the remaining allegations in paragraph 200 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

201. The allegations in paragraph 201 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 201 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

202. The allegations in paragraph 202 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 202 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

203. Wiz denies the allegations in paragraph 203 of the Complaint.

204. Wiz denies the allegations in paragraph 204 of the Complaint.

205. Wiz denies the allegations in paragraph 205 of the Complaint.

206. The allegations in paragraph 206 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 206 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

207. Wiz denies the allegations in paragraph 207 of the Complaint.

208. The allegations in paragraph 208 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 208 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

209. The allegations in paragraph 210 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 210 of the Complaint, and specifically denies that it infringes any valid claim of the '926 patent.

210. Wiz denies the allegations in paragraph 210 of the Complaint.

211. Wiz denies the allegations in paragraph 211 of the Complaint.

212. Wiz denies the allegations in paragraph 212 of the Complaint.

COUNT VI
(INFRINGEMENT OF THE '326 PATENT)

213. Wiz realleges and incorporates by reference its responses to paragraphs 1-35.

214. Wiz admits that what purports to be a copy of U.S. Patent No. 11,775,326 is attached as Exhibit 14 to the Complaint. Wiz further admits that the face of what appears to be the '362 patent indicates that its title is "Techniques for Securing a Plurality of Virtual Machines in a Cloud Computing Environment" and that the date of the patent is October 3, 2023. Wiz denies that the '362 patent was duly and legally issued.

215. Wiz denies that Orca has any right to recover damages for infringement of the '362 patent. Wiz is without knowledge or information sufficient to form a belief as to the truth or falsity of the allegations in paragraph 215 of the Complaint and on that basis, denies them.

216. Wiz denies that the '362 patent is valid and enforceable.

217. Orca's allegations in paragraph 217 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '362 patent.

218. Orca's allegations in paragraph 218 of the Complaint appear intended to reflect the state of the art, claim scope, an appropriate construction of any of the claim terms, the subject matter of the claims, or a claim of infringement, and Wiz therefore denies them. Wiz specifically denies that it infringes any valid claim of the '362 patent.

219. Wiz denies the allegations in paragraph 219 of the Complaint.

220. Wiz admits that paragraph 220 of the Complaint reproduces the language of claim 1 of what appears to be the '362 patent.

221. Wiz denies the allegations in paragraph 221 of the Complaint.

222. Wiz admits the claim language quoted in paragraph 222 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in paragraph 222 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

223. Wiz admits the claim language quoted in paragraph 223 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in

paragraph 223 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

224. Wiz admits the claim language quoted in paragraph 224 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in paragraph 224 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

225. Wiz admits the claim language quoted in paragraph 225 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in paragraph 225 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

226. Wiz admits the claim language quoted in paragraph 226 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in paragraph 226 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

227. Wiz admits the claim language quoted in paragraph 227 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in paragraph 227 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

228. Wiz admits the claim language quoted in paragraph 228 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in paragraph 228 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

229. Wiz admits the claim language quoted in paragraph 229 of the Complaint appears in claim 1 of what appears to be the '362 patent. Wiz denies the remaining allegations in paragraph 229 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

230. Wiz denies the allegations in paragraph 230 of the Complaint.

231. Wiz denies the allegations in paragraph 231 of the Complaint.

232. Wiz denies the allegations in paragraph 232 of the Complaint.

233. Wiz denies the allegations in paragraph 233 of the Complaint.

234. Wiz denies the allegations in paragraph 234 of the Complaint.

235. The allegations in paragraph 235 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 235 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

236. The allegations in paragraph 236 include purported citation to a video posted by Wiz; that video speaks for itself. Wiz denies the remaining allegations in paragraph 236 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

237. The allegations in paragraph 237 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 237 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

238. Wiz denies the allegations in paragraph 238 of the Complaint.

239. Wiz denies the allegations in paragraph 239 of the Complaint.

240. Wiz denies the allegations in paragraph 240 of the Complaint.

241. The allegations in paragraph 241 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 241 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

242. Wiz denies the allegations in paragraph 242 of the Complaint.

243. The allegations in paragraph 243 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 243 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

244. The allegations in paragraph 244 include purported citation to webpages from Wiz's website; that website speaks for itself. Wiz denies the remaining allegations in paragraph 244 of the Complaint, and specifically denies that it infringes any valid claim of the '362 patent.

245. Wiz denies the allegations in paragraph 245 of the Complaint.

246. Wiz denies the allegations in paragraph 246 of the Complaint.

247. Wiz denies the allegations in paragraph 247 of the Complaint.

PRAYER FOR RELIEF

248. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

249. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

250. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

251. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

252. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

253. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

254. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

255. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

256. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

257. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

258. Wiz denies that Orca is entitled to any relief whatsoever, including all relief requested in Orca's "Prayer for Relief." To the extent any statement in the Prayer for Relief is deemed factual and/or requires a response, it is denied.

DEFENSES

259. In addition to denying infringement as to each of Orca's Asserted Patents, and subject to the responses above, Wiz alleges and asserts the following defenses in response to the allegations in Orca's Second Amended Complaint, undertaking the burden of proof only as to those defenses deemed affirmative defenses by law, regardless of how such defenses are denominated herein.

FIRST DEFENSE **(Non-Infringement)**

260. Wiz does not infringe and has not infringed (directly, contributorily, or by inducement), either literally or under doctrine of equivalents, and is not liable for infringement of any valid and enforceable claim of the '031, '032, '685, '809, '926, or '362 patents (collectively, the "Patents-in-Suit")

SECOND DEFENSE **(Invalidity)**

261. The claims of the Patents-in-Suit are invalid and unenforceable under 35 U.S.C § 102 because the claims lack novelty and are taught and suggested by the prior art.

262. The claims of the Patents-in-Suit are invalid and unenforceable under 35 U.S.C § 103 because the claims are obvious in view of the prior art.

263. For example, Orca recently disclaimed every challenged claim of a related patent, U.S. Patent No. 11,431,735, rather than contest that those claims were invalid before the Patent Trial and Appeal Board. *See* Patent Owner's Preliminary Response in *Wiz, Inc. v. Orca Security Ltd.*, Case IPR2024-00220 (April 18, 2024).

264. The claims of the Patents-in-Suit are invalid and unenforceable for failure to satisfy the conditions set forth in 35 U.S.C § 112 including failure to contain a written description, lack of enablement, and indefiniteness because the claims lack novelty and are taught and suggested by the prior art.

THIRD DEFENSE
(Limitations on Patent Damages)

265. Plaintiff's claim for damages, if any, against Wiz for alleged infringement of the Asserted Patents are limited by 35 U.S.C. §§ 286, 287, and/or 288.

FOURTH DEFENSE
(Prosecution History Estoppel)

266. By reason of statements, representations, concessions, admissions, arguments, and/or amendments, whether explicit or implicit, made by or on behalf of the applicant during the prosecution of the patent applications that led to the issuance of the Asserted Patents, Plaintiff's claims of infringement are barred in whole or in part by the doctrine of prosecution history estoppel.

FIFTH DEFENSE
(Patent Marking)

267. Any claim for damages for patent infringement is limited by 35 U.S.C. § 287 to those damages occurring only after the notice of infringement.

SIXTH DEFENSE
(License)

268. Plaintiff's claims are barred in whole or in part by an express or implied license and/or the patent exhaustion doctrine.

SEVENTH DEFENSE
(Non-Exceptional Case)

269. Plaintiff cannot prove that this is an exceptional case that would justify an award of attorney's fees against Wiz pursuant to 35 U.S.C. § 285.

EIGHTH DEFENSE
(Ensnarement)

270. Plaintiff's claims for infringement are barred by the doctrine of ensnarement.

NINTH DEFENSE
(Unclean Hands)

271. Plaintiff's claims are barred by the doctrine of unclean hands. For example, Orca has obtained non-public, proprietary information of Wiz without Wiz's authorization on multiple occasions and, rather than return or destroy that information, has used that information to bring baseless litigation targeting Wiz, including this case.

272. As a specific example, in this action, Exhibit 4 to Orca's original complaint (D.I. 1-1 at 45-55) is from a non-publicly accessible webpage behind a login page specifically designed to prevent access to Wiz's proprietary information unless authorized by Wiz. When Wiz confronted Orca with this information, Orca stated it obtained this document "from a third party some time in 2022." Notably, Orca continued to amend claims of the Asserted Patents in 2022 and later.

273. Similarly, in August 2022, Orca filed a complaint against one of Wiz's employees and sought to prevent her from working for Wiz for a year, accusing her of misusing Orca's confidential information based only on the fact that she had previously worked for Orca and was now working for Wiz. *See Orca Security, Inc. v. Jacques*, 1:22-cv-01048-CFC (D. Del. August 10, 2022), D.I. 10. However, Orca's own briefing showed it was **Orca** that improperly handled Wiz's confidential information. In its submissions in that case, Orca acknowledged that in the industry, "Request For Proposal ('RFP') processes" involve "blind proposals, including confidential price and technological information, such that competitors like Orca and Wiz do not

know what is contained in each other's proposal or even which competitors they are up against[,]” but Orca brazenly admitted that a potential client had accidentally given Orca documents relating to Wiz's response to an RFP and—instead of promptly returning or destroying them—Orca used them as the purported basis for the suit. *Id.* at 4, 8. On information and belief, Orca was also asked to delete the information. Orca voluntarily dismissed that suit before Wiz's employee filed any response. Again, Orca continued to amend claims of the Asserted Patents after August 2022.

TENTH DEFENSE
(Equitable Estoppel/Waiver)

274. Orca's claims are barred in whole or in part by the doctrines of equitable estoppel or waiver.

RESERVATION OF RIGHTS

275. Wiz reserves the right to amend this Answer to Orca's Second Amended Complaint and assert further affirmative defenses in the event that discovery indicates that doing so would be appropriate.

PRAYER FOR RELIEF

WHEREFORE, Wiz respectfully requests that the Court enter judgment in favor and against Plaintiff as follows:

1. Dismissing with prejudice Plaintiff's claims against Wiz;
2. Denying all relief that Plaintiff seeks in its Complaint;
3. Finding this case exceptional under 35 U.S.C. § 285 and awarding Wiz all costs and attorney's fees; and
4. Awarding any other relief the Court deems just and equitable.

DEMAND FOR A JURY TRIAL

In accordance with Fed. R. Civ. P. 38, Wiz demands a trial by jury on all issues so triable.

WIZ'S COUNTERCLAIMS

Counterclaim-Plaintiff Wiz hereby alleges the following Counterclaims against Counterclaim-Defendant Orca:

Wiz is a Technology Success Story

1. Wiz is a technology success story and one of the hottest startups in the world. Wiz's products secure their customers' use of the "cloud"—*i.e.* the now ubiquitous software running on servers provided by Amazon Web Services, Google Cloud, Microsoft Azure and others. With cloud computing, a user does not need to have a local or personal computer capable of doing all of their tasks; instead, a computing "workload" can be hosted on remote servers in the "cloud." With the incredible importance of the "cloud," Wiz has provided immense value by securing and protecting these assets.

2. Though founded only in 2020, Wiz currently has a \$12 billion dollar valuation, and recorded over \$350 million in sales in annual recurring revenue. Its customers are a "who's who" of private industry, including Morgan Stanley, BMW, DocuSign, Salesforce, Fox Corporation, Colgate-Palmolive, among many others. More than 40% of the Fortune 100 are Wiz customers.

3. Wiz's success was not built overnight, however. Wiz was founded in 2020 by serial cybersecurity entrepreneurs with over a decade of expertise in the industry. The founders of Wiz, Assaf Rappaport, Yinon Costica, Roy Reznik, and Ami Luttwak served in the Israeli Defense Forces ("IDF"), including various members in Unit 8200, an elite intelligence division and Unit 81, a secret technology section part of the Special Operations Division of the Military Intelligence Directorate of the IDF.

4. In 2012, Rappaport, Luttwak, and Reznik created Adallom, a cloud security company that was based in Menlo Park, California. Ahead of its time, it secured companies' use

of enterprise software cloud applications, such as SharePoint, Dropbox and Box. After raising tens of millions of dollars in venture capital funding, Adallom was acquired by Microsoft for approximately \$320 million in 2015. Microsoft then turned to Rappaport to lead its new cloud security division, with Rappaport later operating as general manager for Microsoft's entire research and development center in Israel.

5. After five years at Microsoft running its cloud security division, Rappaport and the future Wiz founders left Microsoft to build a new venture. The entrepreneurship bug continued to push them, and many investors kept encouraging them to start a new venture.

6. Wiz was incorporated in January 2020 and exited stealth in December of that year, when it announced a \$100 million Series A funding.³ Backed by seed funding from venture capital, the Wiz founders considered a few different ideas. But after talking to dozens of prospective buyers about what they needed most, the group quickly focused on cloud visibility.

7. Wiz's products and services have been a hit, generating hundreds of millions of dollars in revenue. Wiz has also made repeated headlines by publicly identifying cloud security vulnerabilities in the industry. *See* Ex. F, Microsoft *Patched Bing Vulnerability That Allowed Snooping on Email and Other Data*, WALL STREET JOURNAL, March 29, 2023, ("The problem was discovered by outside researchers at the security firm Wiz Inc.").

³ <https://www.wiz.io/blog/wiz-comes-out-of-stealth-with-100m-series-a-funding-to-reinvent-cloud-security>

Wiz's Innovative Intellectual Property Portfolio

8. Wiz's investment in research and development of new cloud security technology and features has led to Wiz being one of the most innovative cybersecurity companies in the industry. This is also reflected in Wiz's intellectual property portfolio.

9. Some of the advancements developed by Wiz related to developing a new holistic cloud security solution that, instead of focusing on each "workload" within the cloud to detect vulnerabilities, focuses on providing visibility across the entire cloud environment to provide unified risk analysis and address risk across a user's entire deployment. These solutions include those directed to, among other things, creating a network graph that visually represents network objects for the user, generating unified graph models across multiple cloud computing platforms by genericizing and imputing network entities, and applying policies on a path through the network to mitigate risks. As claimed in Wiz's patents, these were insights by Wiz to change the model for cloud security, driven by thinking about visibility across the entire cloud, rather than focusing on each cloud server or "workload." The marketplace has agreed.

10. Wiz has also innovated in the area of Artificial Intelligence ("AI") and cloud security. For example, Wiz has developed solutions directed to, among other things, detecting cybersecurity risks from AI models and utilizing Large Language Models ("LLMs") to assist in responding to cybersecurity incidents.

11. Wiz applied for and received patent protection from the United States Patent and Trademark Office ("USPTO") for these advancements, including patents infringed by Orca as discussed further below.

Orca’s Lack of Success in the Marketplace, Followed by Copying of Wiz

12. In contrast to Wiz, Orca has lagged in developing features and in the marketplace. Founded in 2019—only a year prior to Wiz—Orca markets itself as a cloud cybersecurity company. Orca originally focused on specific legacy risks such as workload security. On information and belief, Orca’s purportedly new agentless approach to workload scanning, which it would come to call “SideScanning” (*see, e.g.*, D.I. 1-1 at Ex. 3), did not originally include Wiz’s patented features. For example, according to Orca’s own early descriptions, its “SideScanning” approach works on a per-workload basis. With SideScanning, Orca collects workload data and “then reconstructs the workload’s file system – OS, applications, and data – in a virtual read-only view” to perform its risk analysis.⁴ Orca continues to tout “agentless SideScanning” as its primary “innovation” today.⁵ However, agentless security solutions for the cloud have been known in the industry since before Orca existed, and many in the industry use agentless techniques today,⁶ including the cloud service providers themselves.⁷

13. Orca has not experienced as much success with its approach. While the company claimed it expected “triple figure growth” in 2023, it reportedly laid off 15% of its workforce in

⁴ <https://orca.security/platform/agentless-sidescanning/>

⁵ <https://orca.security/about/>

⁶ *See, e.g.*, <https://blogs.cisco.com/security/agentless-threat-detection-for-microsoft-azure-workloads-with-cisco-stealthwatch-cloud>; <https://www.paloaltonetworks.com/cyberpedia/what-is-the-difference-between-agent-based-and-agentless-security>

⁷ *See, e.g.*, <https://aws.amazon.com/about-aws/whats-new/2023/11/amazon-inspector-agentless-assessments-ec2-preview/>; <https://learn.microsoft.com/en-us/azure/defender-for-cloud/concept-agentless-data-collection>

January 2024. It publicly stated that the layoffs were due to “current macroeconomic conditions.”⁸

14. In its efforts to catch up to Wiz, Orca has repeatedly reviewed and kept Wiz confidential materials meant for potential and current customers, including attaching one such document to its complaint against Wiz. Exhibit 4 to Orca’s complaint is a Wiz document that is not public and only accessible through a login page specifically designed to prevent access to Wiz’s proprietary information without Wiz’s authorization. When Wiz asked Orca how it obtained this document, Orca would only respond that it received it “from a third party some time in 2022.”


15. As another example, in August 2022, Orca filed a complaint against one of Wiz’s employees, accusing her of misusing Orca’s confidential information because she changed jobs, but Orca’s own briefing showed it was Orca—not Wiz or its employee—that had improperly handled its competitor’s confidential information. *See Orca Security, Inc. v. Jacques*, 1:22-cv-01048-CFC (D. Del. August 10, 2022), D.I. 10. Orca brazenly admitted that a potential client had accidentally given Orca documents relating to Wiz’s response to a Request for Proposal (“RFP”) and—instead of promptly returning or destroying them—Orca kept those documents even though it knew such RFPs are supposed to be “blind” and include “confidential price and technological information, such that competitors like Orca and Wiz do not know what is contained in each other’s proposal or even which competitors they are up against[.]” *Id.* at 4, 8. Orca voluntarily dismissed that suit before Wiz’s employee responded.

⁸ See <https://www.calcalistech.com/ctechnews/article/skm23oz00t>; <https://www.crn.com/news/security/2024/orca-security-cuts-15-percent-of-staff>


16. After Wiz showed its successful approach in the market, Orca has tried to compete by repeatedly copying features first debuted by Wiz. This includes adopting features patented by Wiz as discussed further below.

17. Rather than “copying” the idea of serving coffee at a conference, Orca’s copying has included adopting Wiz’s patented technology. Orca regularly copies Wiz’s features after they are released:

18. In June 2022, Wiz announced its new Cloud Detection and Response (CDR) feature, which provided capabilities to simulate, detect, and respond to cloud security events. See <https://www.wiz.io/blog/uniting-builders-and-defenders-a-new-vision-for-cloud-security>. A month later, in July 2022, Orca announced a feature with the same name. See <https://orca.security/resources/blog/orca-cloud-security-platform-adds-cloud-detection-and-response/>.



19. In December 2021, Wiz announced it was offering tools for its customers to operationalize a Shift-Left strategy. See <https://www.wiz.io/blog/wiz-magic-shifts-left> (dated December 9, 2021). Approximately six months later, Orca announced its was adding the same features. See <https://orca.security/resources/blog/shift-left-security-platform/> (dated May 11, 2022).



20. Wiz has provided Cloud Infrastructure Entitlement Management or “CIEM” since its earliest product offerings, including using those terms long before Orca. On February 10, 2022, Orca released its CIEM feature. See <https://orca.security/resources/press-releases/orca-platform-expanded-ciem-multi-cloud-security-score> (dated Feb. 10, 2022).

21. By at least August 2021, Wiz had released its Threat Center, a dedicated feed of high-profile security issues with analysis of their impact on your organization. See <https://www.wiz.io/blog/protecting-your-environment-from-chaosdb> (dated Aug. 29, 2021, stating “Wiz Threat Center shows you which assets are at-risk to the most dangerous threats”).

In March 2022, Orca released its version of the same feature, redesigning its interface to be a “news feed” and including a “From the news” functionality similar to Wiz. *See*

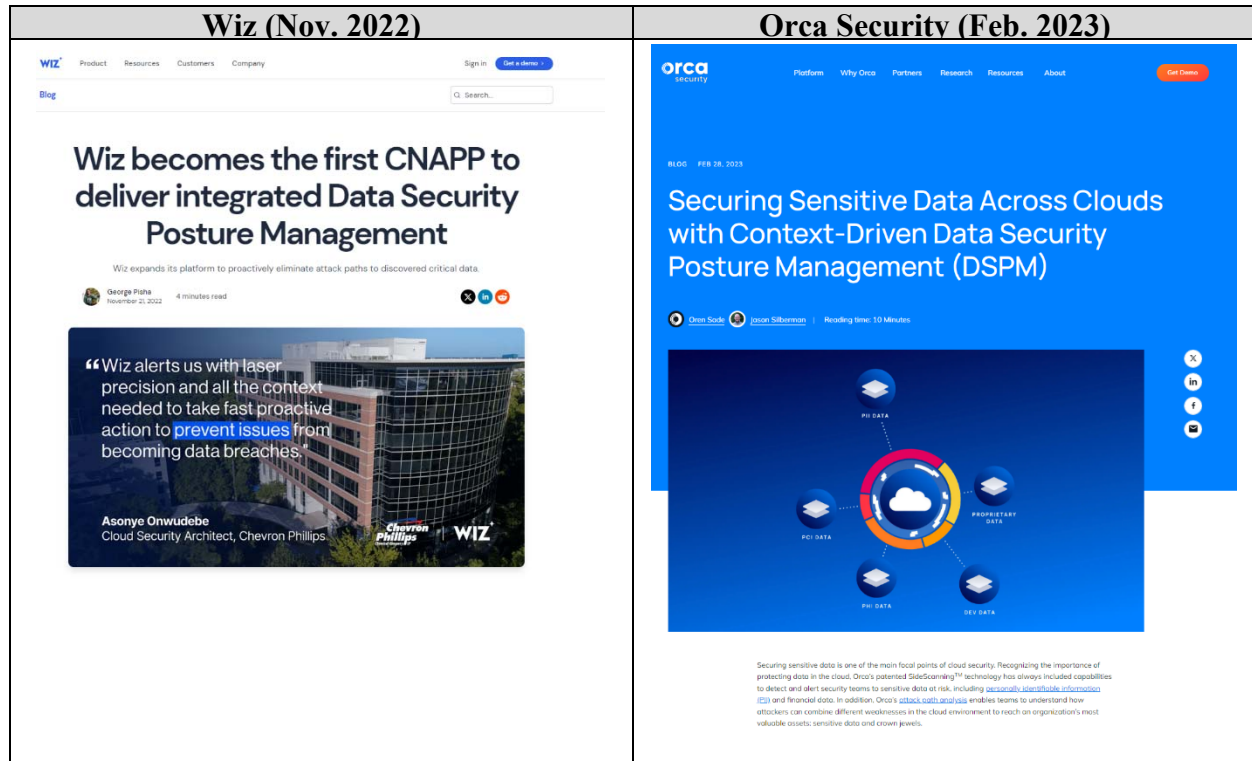
<https://orca.security/resources/blog/enhancing-the-orca-security-risk-dashboard-with-an-integrated-news-feed/> (dated Mar. 8, 2022).

22. On November 21, 2022, Wiz announced it was the first cloud security platform to offer Data Security Posture Management (DSPM) capabilities to continuously monitor for critical data exposure so organizations can respond before a breach occurs. *See*

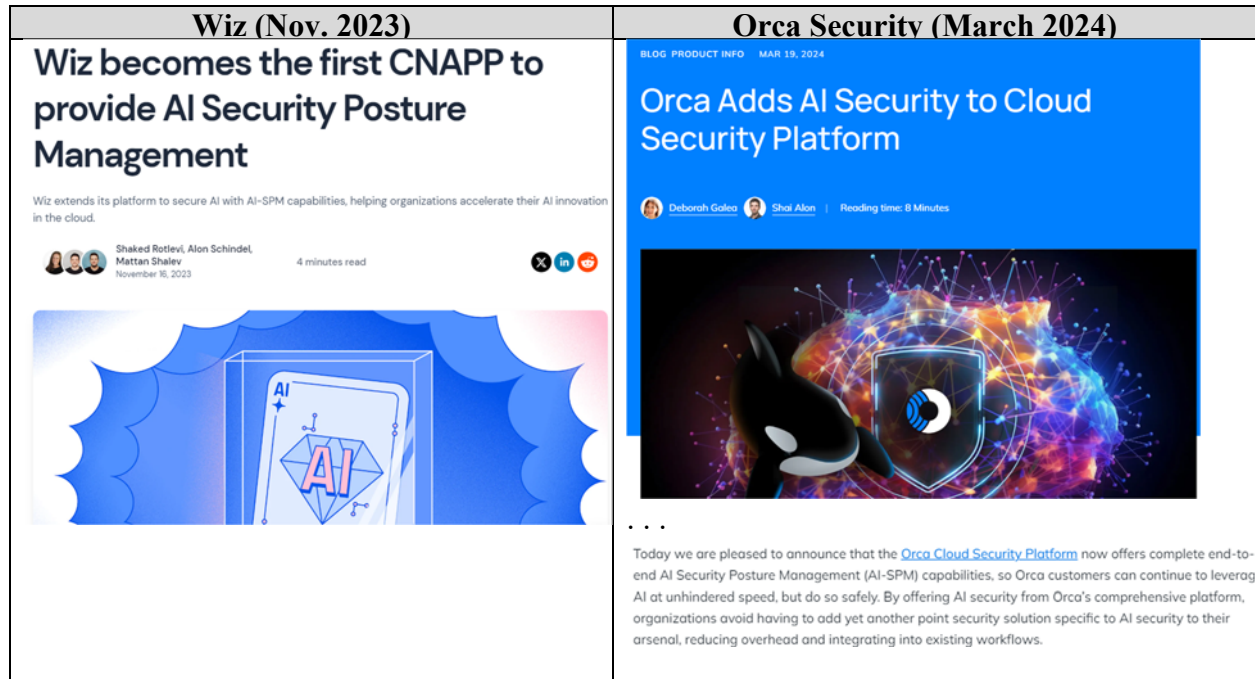
<https://www.wiz.io/blog/wiz-becomes-first-cnapp-to-deliver-integrated-data-security-posture-management> (dated Nov. 21, 2022). Months later, on Feb. 28, 2023, Orca announced it was

launching the same feature. *See* <https://orca.security/resources/blog/securing-sensitive-data-across-clouds-with-data-security-posture-management-dspm/> (blog post dated Feb. 28, 2023

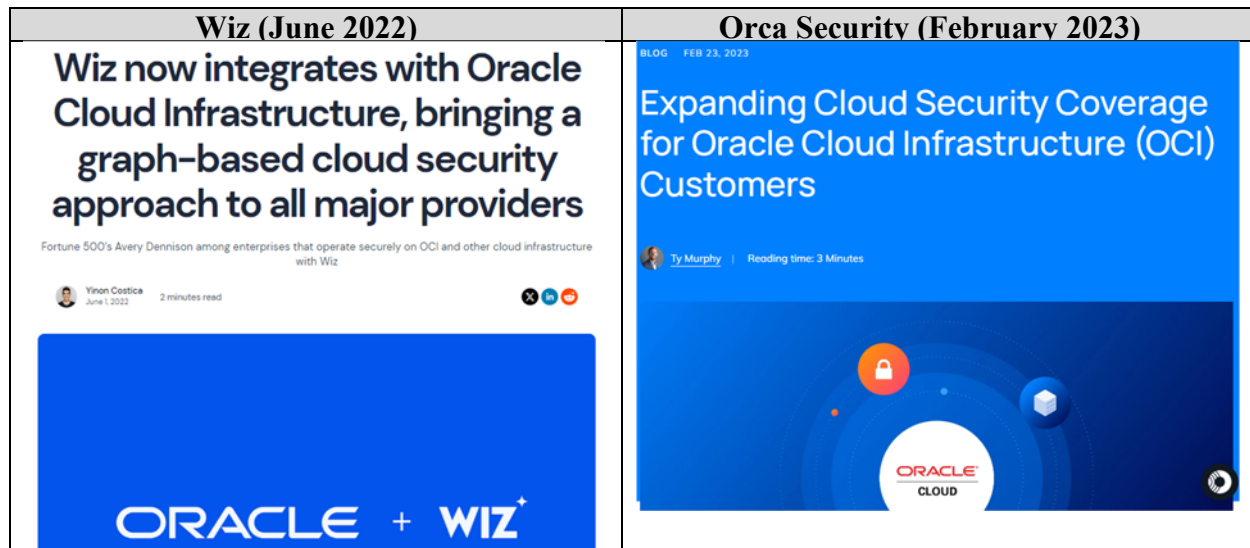
stating, in part: “Today, we are excited to announce that we have now significantly expanded our cloud data security coverage and capabilities, launching a comprehensive offering of Data Security Posture Management (DSPM) as part of the Orca Cloud Security Platform.”).



23. On November 16, 2023, Wiz announced it was the first cloud security platform to secure AI with AI Security Posture Management or “AI-SPM capabilities”—a term coined by Wiz—including by announcing an “AI Security Dashboard.” See <https://www.wiz.io/blog/ai-security-posture-management> (dated November 16, 2023). Four months later, on March 19, 2024, Orca announced it would also be offering “AI Security Posture Management (AI-SPM) capabilities”—using the same term coined by Wiz—and an “AI Security dashboard.” See <https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/> (dated March 19, 2024).



24. On June 1, 2022, Wiz announced that its platform could integrate with the Oracle Cloud Infrastructure (OCI). See <https://www.wiz.io/blog/supporting-oracle-cloud-wiz-brings-the-first-graph-based-cloud-security-approach-to-all-major-providers> (dated June 1, 2022). On February 23, 2023, Orca announced its platform could now connect with OCI. See <https://orca.security/resources/blog/expanding-cloud-security-coverage-for-oracle-cloud-infrastructure/> (dated Feb. 23, 2023).



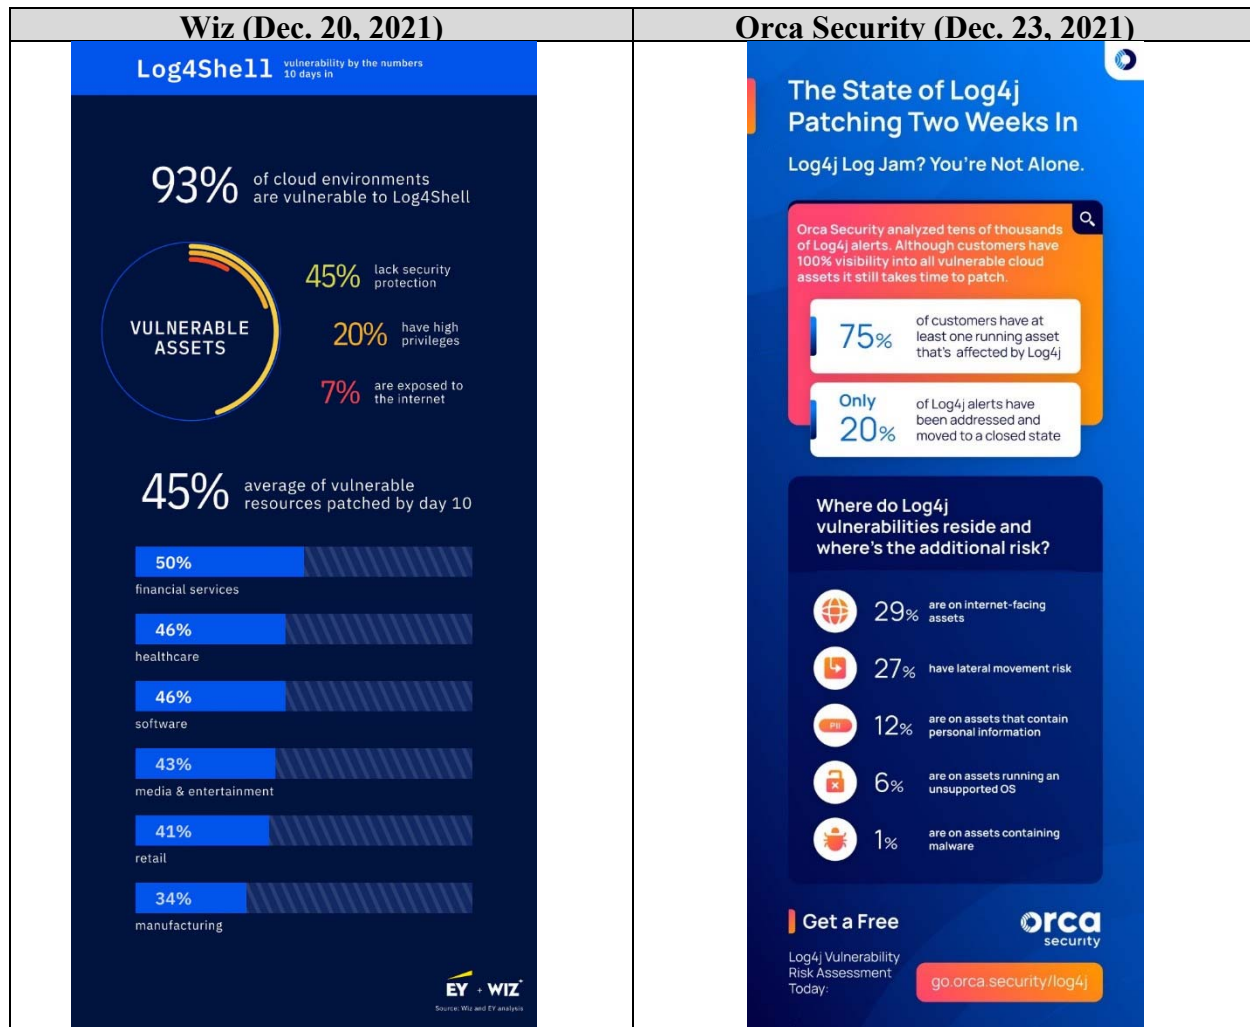
25. Similarly, Wiz announced support for Alibaba Cloud on June 29, 2022. *See* <https://www.wiz.io/blog/wiz-extends-cnapp-leadership-with-protection-for-alibaba-cloud> (dated June 29, 2022). On November 9, 2022, Orca announced its platform would also support Alibaba Cloud. *See* <https://orca.security/resources/blog/expanding-cloud-security-for-alibaba-cloud/> (dated Nov. 9, 2022).



26. On June 21, 2023, Wiz announced it became the first CNAPP to provide an end-to-end cloud forensics experience. *See* <https://www.wiz.io/blog/wiz-becomes-the-first-cnapp-to-provide-end-to-end-cloud-forensics-experience> (dated June 21, 2023). On April 29, 2024, Orca announced it was launching a cloud digital forensics and incident response service. *See*

<https://orca.security/resources/press-releases/orca-security-launches-cloud-digital-forensics-and-incident-response-service-to-empower-rapid-investigation-of-cloud-incidents/> (dated April 29, 2024).

27. Orca's copying, however, has included the mundane as well, including copying infographics from Wiz. In the wake of the Log4Shell vulnerability that was major news in the cybersecurity industry, Wiz published a statistics blog about the prevalence of the issue, including an infographic. See <https://www.wiz.io/blog/10-days-later-enterprises-halfway-through-patching-log4shell> (dated Dec. 20, 2021). Three days later, Orca published a blog post with a very similar graphic. See <https://orca.security/resources/blog/instantly-detect-log4j-vulnerabilities-on-aws-azure-and-google-cloud/> (dated Dec. 23, 2021).



28. Based on the above examples, Orca appears to have a culture of copying Wiz's innovations, including its patented technology as shown herein.

29. Wiz did not choose to bring this litigation, but faced with Orca's meritless claims, it is now forced to correct the record about Wiz's innovation, Orca's copying of Wiz, and Orca's use of Wiz's intellectual property. Orca is improperly using Wiz's inventions, specifically those claimed in U.S. Patent Nos. 11,722,554 (the "'554 Patent"); 11,929,896 (the "'896 Patent"); 11,936,693 (the "'693 Patent"); 12,001,549 (the "'549 Patent"); and 12,003,529 (the "'529 Patent") (collectively, Wiz's "Asserted Patents"), as discussed in further detail below. Exs. A-E. Wiz brings these counterclaims to address that infringement.

30. On information and belief, Orca was aware the Asserted Patents that issued prior to June 4, 2024, and Orca's infringement thereof prior to these counterclaims at least due to Orca's monitoring of Wiz patents. As one example, in its original complaint in this action, Orca cited and quoted from Wiz's U.S. Patent No. 11,374,982, demonstrating its monitoring of Wiz's patent portfolio. *See, e.g.*, D.I. 1, ¶ 22. Moreover, Orca has a demonstrated culture of copying Wiz, as discussed above. In addition, Orca is aware of each of the Asserted Patents and its infringement thereof at least as of the filing of these counterclaims. Orca has and continues to willfully infringe all of the Asserted Patents.

THE PARTIES

31. Wiz is a Delaware corporation with its principal place of business at One Manhattan West, 57th Floor, New York, New York.

32. On information and belief, Orca is an Israeli company with its principal place of business at 3 Tushia St., Tel Aviv, Israel 6721803.

JURISDICTION AND VENUE

33. This Court has subject matter jurisdiction over the matters asserted herein under 28 U.S.C. §§ 1331 and 1338(a).

34. Orca is subject to this Court's personal jurisdiction at least because Orca initiated this lawsuit and, on information and belief, Orca's business operations include software and services for use in Delaware.

35. Venue is proper in this District pursuant to 28 U.S.C. Section 1391 (b), (c), and/or 1400(b), at least because Orca is a foreign entity that, on information and belief, has committed acts of infringement in this District.

COUNTERCLAIM I
(INFRINGEMENT OF U.S. PATENT NO. 11,722,554)

36. Wiz is the sole and exclusive owner, by assignment, of all rights, title and interest in U.S. Patent No. 11,722,554 (the “’554 patent”), entitled “System and method for analyzing network objects in a cloud environment.” The ’554 patent was duly and legally issued by the U.S. Patent and Trademark Office on August 8, 2023. The named inventors of the ’554 patent are Shai Keren, Danny Shemesh, Roy Reznik, Ami Luttwak, and Avihai Berkovitz. A copy of the ’554 patent is attached as Exhibit A.

37. The ’554 patent generally relates to determining abnormal configuration of network objects in a cloud environment for security purposes. *See* ’554 patent at 2:51-67. This is done through the use of a network graph that includes a visual representation of network objects in the cloud computing environment. Describing an embodiment shown in figure 2 the patent states, “[a] network graph is a data feature describing the various objects included in, and adjacent to, a network, as well as the relationship between such objects.” *Id.* at 8:49-51. The network graph includes network object relationships. “Network object relationships are descriptions of the various connections between the network objects identified at S210. Network relationships may describe aspects of the connections between objects including, without limitation, connected objects, relevant ports of connected objects, connection bandwidths, connection durations, connection protocols, connection names or IDs, connection statuses, and the like, as well as any combination thereof.” *Id.* at 9:4-11.

38. The ’554 patent discloses using the network graph and network object relationships to generate at least one network insight. *See Id.* at 2:63-67. The network insights “are natural-language representations of aspects of the network graph[.]” *Id.* at 10:15-17. “Network insights may include a pure-text descriptions of objects and relationships.” *Id.* at 10:17-

19. “In addition, network insights may include detailed descriptions of objects, relationships and the like as well as any combination thereof.” *Id.* at 10:25-27.

39. Orca has infringed and continues to directly infringe one or more claims of the ’554 patent by making, using, selling, offering for sale, and/or importing into the United States without authority or license, the Orca Platform with Attack Path Analysis in violation of 35 U.S.C. § 271(a). Orca’s infringement includes infringement of, for example, claim 1 of the ’554 patent.

40. Claim 1 of the ’554 patent recites:

1. A method for determining abnormal configuration of network objects deployed in a cloud computing environment, comprising:
collecting network object data on a plurality of network objects deployed in the cloud computing environment;
constructing a network graph based on the collected network object data, wherein the network graph includes a visual representation of network objects identified in the cloud computing environment;
determining relationships between the identified network objects in the network graph, wherein the determined relationships between the identified network objects includes descriptions of connections between the identified network objects;
analyzing the network graph and the determined relationships to generate insights, wherein the generated insights include at least a list of abnormal connections between the identified network objects; and
tagging network objects in the network graph for which the insight is generated.

41. On information and belief, Orca practices each and every limitation of claim 1 of the ’554 patent by and through the use of the Attack Path Analysis.

42. The preamble of claim 1 recites “[a] method for determining abnormal configuration of network objects deployed in a cloud computing environment, comprising:” To the extent the preamble is limiting, Orca practices this step by, for example, using its Attack Path Analysis product to collect data on assets in cloud computing environments. *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder,*

<https://orca.security/resources/blog/cloud-attack-path-analysis/> (“Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.”).

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the core capabilities of a Cloud-Native Application Protection Platform (CNAPP):

43. Claim 1 further recites “collecting network object data on a plurality of network objects deployed in the cloud computing environment” Orca’s public blog posts confirm that Orca practices this step by, for example but not limited to, Orca’s Attack Path Analysis “collects and correlates contextual data on each asset” and that “it is important to view risks as an interrelated chain, rather than just siloed risks.” *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/> (“To fully understand where your organization’s most critical weaknesses are, it is important to view risks

as an interrelated chain, rather than just siloed risks. By understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard.”).

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

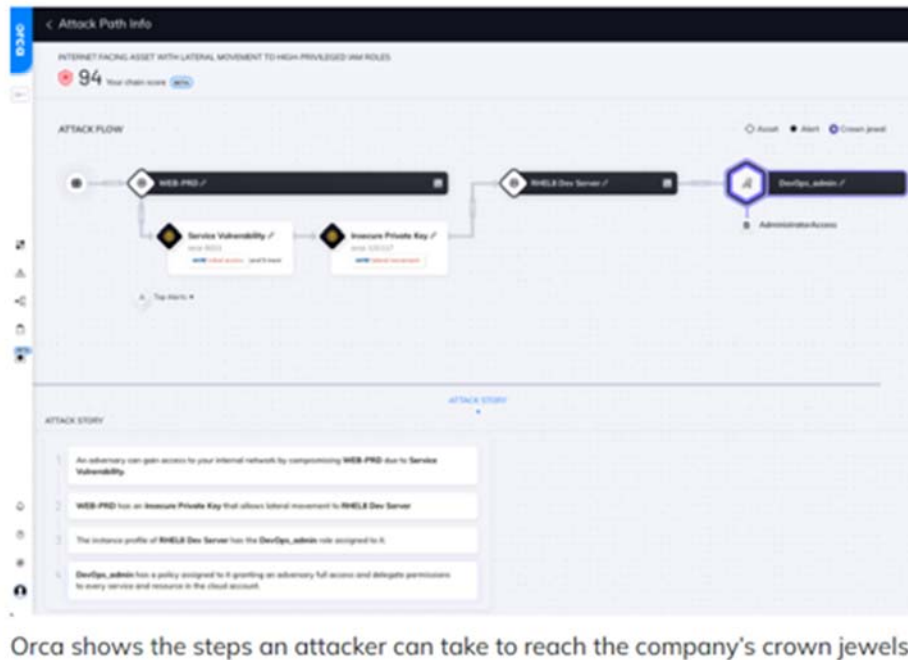
Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization’s most valuable assets, or ‘crown jewels’, and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates [contextual data](#) on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the [core capabilities of a Cloud-Native Application Protection Platform \(CNAPP\)](#):

To fully understand where your organization's most critical weaknesses are, it is important to view risks as an interrelated chain, rather than just siloed risks. By understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard.



See also, *Data Security and Posture Management*, <https://orca.security/platform/data-security-and-posture-management-dspm/> (“Orca’s DSPM dashboard provides data security teams with an overview of their cloud data stores, sensitive data, and security and compliance alerts. Orca scans managed, unmanaged, and shadow data, giving security teams wide and deep visibility

into where their data resides.”).

Discover and classify data in your cloud

Orca's DSPM dashboard provides data security teams with an overview of their cloud data stores, [sensitive data](#), and security and compliance alerts. Orca scans managed, unmanaged, and shadow data, giving security teams wide and deep visibility into where their data resides.

- ✓ Get a multi-cloud inventory of data storage assets—including databases, and files in virtual machines, storage buckets, and containers.
- ✓ Know which data stores contain sensitive data and of what type—including PII, [PCI](#), PHI, or financial information—for both security and regulatory purposes.
- ✓ Leverage interactive graphs that show the location and relationship between data stores and other cloud entities.

44. Claim 1 further recites “constructing a network graph based on the collected network object data, wherein the network graph includes a visual representation of network objects identified in the cloud computing environment” Orca’s public blog posts confirm that Orca practices this step by, for example but not limited to, describing operation of Orca’s Attack Path Analysis as “representing attack paths in a visual graph with contextual data on all relevant cloud entities” and by using their “attack path visualization” tool. *See, e.g., Cloud*

Attack Path Analysis: Work Smarter Not Harder, <https://orca.security/resources/blog/cloud-attack-path-analysis/>

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates [contextual data](#) on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the [core capabilities of a Cloud-Native Application Protection Platform \(CNAPP\)](#):

45. Claim 1 further recites “determining relationships between the identified network objects in the network graph, wherein the determined relationships between the identified network objects includes descriptions of connections between the identified network objects” Orca’s public blog posts and marketing videos confirm that Orca practices this step by, for example but not limited to, using Orca’s Attack Path Analysis to identify “risk combinations” between objects including “vulnerability status, misconfiguration risks, trust and authorization[.]” *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*,

<https://orca.security/resources/blog/cloud-attack-path-analysis/>.

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

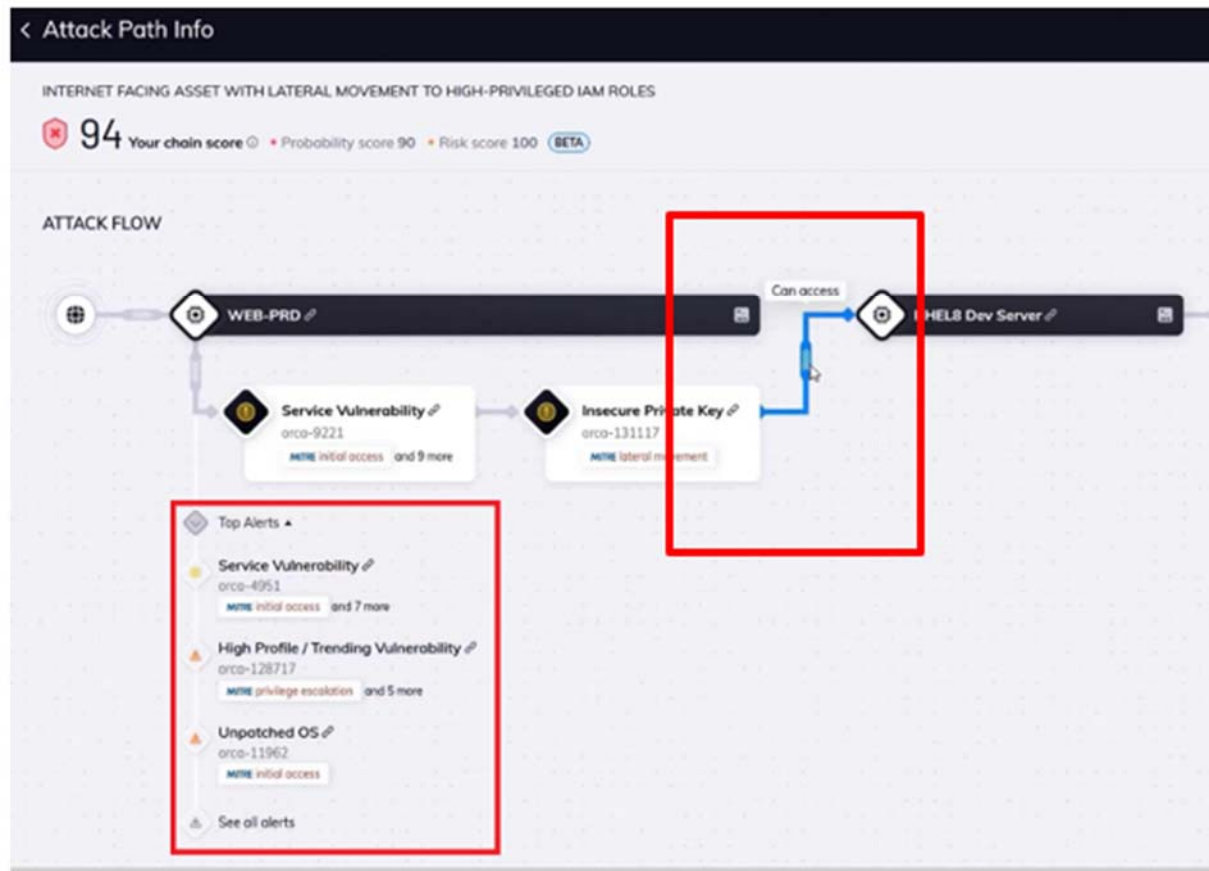
Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates [contextual data](#) on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the [core capabilities of a Cloud-Native Application Protection Platform \(CNAPP\)](#):

See, also, *Orca Bytes: Attack Path Analysis*, <https://www.youtube.com/watch?v=MJkO8UfQa-8>.

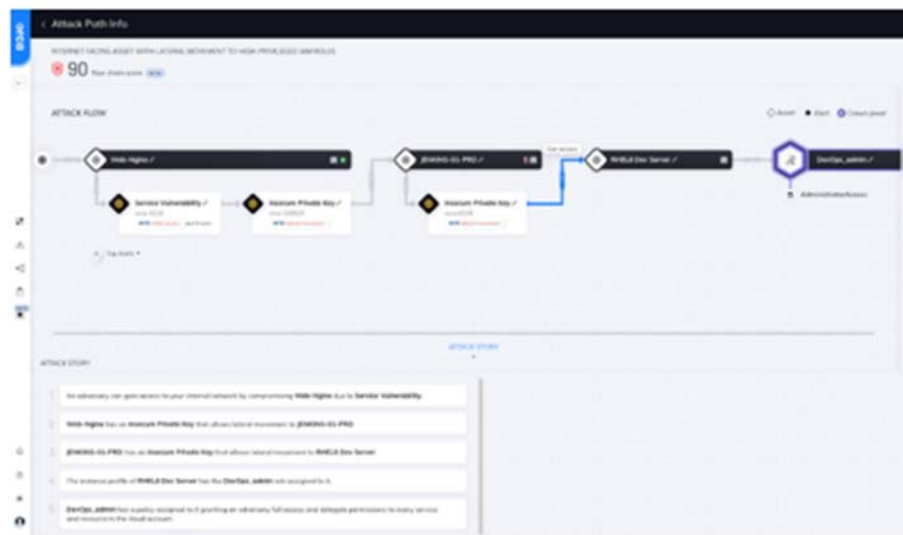


46. Claim 1 further recites “analyzing the network graph and the determined relationships to generate insights, wherein the generated insights include at least a list of abnormal connections between the identified network objects” Orca’s public blog posts and marketing videos confirm that Orca practices this step by, for example but not limited to, using Orca’s Attack Path Analysis’s “analysis and prioritization capabilities” and “representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.” See, e.g., *Cloud Attack Path Analysis: Work Smarter Not Harder*,

<https://orca.security/resources/blog/cloud-attack-path-analysis/>.

Prioritize Attack Paths, Not Just Siloed Alerts

With Orca's new attack path analysis and prioritization capabilities, security teams no longer need to sift through hundreds of siloed alerts to find out which issues need to be fixed and in what order, but instead can now focus on a much smaller amount of prioritized attack paths, or combinations of alerts, that endanger the company's most critical assets. By prioritizing and scoring each attack path, Orca pinpoints exactly which risks need to be remediated to 'break the chain'.



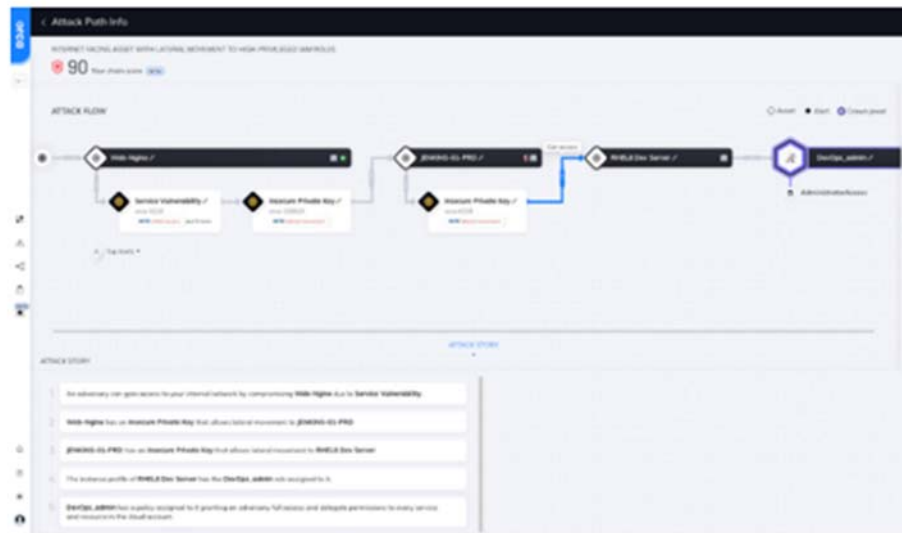
Attack path visualization of how an attacker can access sensitive company data

47. Finally, Claim 1 further recites “and tagging network objects in the network graph for which the insight is generated.” Orca’s public blog posts and marketing videos confirm that Orca practices this step by, for example but not limited to, using its Attack Path Analysis for “prioritizing and scoring each attack path[.]” See, e.g., *Cloud Attack Path Analysis: Work*

Smarter Not Harder, <https://orca.security/resources/blog/cloud-attack-path-analysis/>.


Prioritize Attack Paths, Not Just Siloed Alerts

With Orca's new attack path analysis and prioritization capabilities, security teams no longer need to sift through hundreds of siloed alerts to find out which issues need to be fixed and in what order, but instead can now focus on a much smaller amount of prioritized attack paths, or combinations of alerts, that endanger the company's most critical assets. By prioritizing and scoring each attack path, Orca pinpoints exactly which risks need to be remediated to 'break the chain'.



Attack path visualization of how an attacker can access sensitive company data

See, e.g., *Orca Bytes: Attack Path Analysis*, <https://www.youtube.com/watch?v=MJkO8UfQa-8>.



48. On information and belief, Orca was aware of the '554 patent and Orca's infringement thereof prior to these counterclaims at least due to Orca's monitoring of Wiz patents as shown by, in its original complaint in this action, that Orca cited and quoted from Wiz's U.S. Patent No. 11,374,982. *See, e.g.*, D.I. 1, ¶ 22. Further, as demonstrated above Orca has repeatedly shown a culture of copying Wiz. This is just one more example of Orca seeing Wiz's success and copying instead of innovating. Moreover, Orca is aware of the '554 patent and Orca's infringement thereof at least as of the filing of these counterclaims. Accordingly, Orca has and continues to willfully infringe the '554 patent.

49. Orca has induced and continues to induce infringement of one or more claims of the '554 patent by, for example but not limited to, encouraging customers to use its Attack Path

Analysis in a manner that directly infringes those claims. Despite its knowledge of the existence of the '554 patent, since at least the filing of this Counterclaim, Orca, upon information and belief, continues to encourage, instruct, enable and otherwise cause its customers to use its Attack Path Analysis in a manner that infringes one or more claims of the '554 patent. Upon information and belief, Orca specifically intends that its customers use its Attack Path Analysis in a manner that infringes one or more claims of the '554 patent by, at a minimum, providing instructions and/or support documentation directing customers on how to use its Attack Path Analysis in an infringing manner, in violation of 35 U.S.C. § 271(b). For example, Orca's public blog posts cited above provide instructions and encourage customers to practice all steps of the claimed method stating: "To fully understand where your organization's most critical weaknesses are, it is important to view risks as an interrelated chain, rather than just siloed risks. By understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard." *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/>. Further, Orca provides video explanation of Attack Path Analysis stating, "with Orca's new attack path analysis and prioritization capabilities, security teams can now laser focus on a small number of prioritized attack paths or alert on combinations that endanger the company's most critical assets, and every path and link in the path is scored so you can pinpoint exactly which risks need to be remediated." *See Orca Bytes: Attack Path Analysis* (Mar. 31, 2022), <https://www.youtube.com/watch?v=MJkO8UfQa-8>.

50. Orca has contributed and continues to contribute to the infringement of one or more claims of the '554 patent. Upon information and belief, Orca knows that its Attack Path

Analysis feature is especially made and/or adapted for users to infringe one or more claims of the '554 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use. Because Orca included the features, such as for example but not limited to Attack Path Analysis, in Orca's products, Orca intends for customers to use it. Upon information and belief, its Attack Path Analysis feature has no suitable use that is non-infringing, and therefore Orca intends for customers to use its Attack Path Analysis in an infringing manner. Orca's sales of products including its Attack Path Analysis constitute contributory infringement in violation of 35 U.S.C. § 271(c).

COUNTERCLAIM II
(INFRINGEMENT OF U.S. PATENT NO. 11,929,896)

51. Wiz is the sole and exclusive owner, by assignment, of all rights, title and interest in U.S. Patent No. 11,929,896 (the "'896 patent"), entitled "System and Method for Generation of Unified Graph Models for Network Entities." The '896 patent was duly and legally issued by the U.S. Patent and Trademark Office on Mar. 12, 2024. The named inventors of the '896 patent are Daniel Shemesh, Liran Moysi, Roy Reznik, and Shai Keren. A copy of the '896 patent is attached as Exhibit B.

52. The '896 patent generally relates to generation of network graph models for network entities. *See* '896 patent at 2:48-59. This is done by collecting network entities and network entity properties, and creating a network graph that is a multi-dimensional data structure representing the network entities. *Id.* Describing an embodiment shown in Figure 2 the patent states, a network graph is generated. A graph is a multi-dimensional data feature providing a representation of the contents and structure of a network, cloud, environment, or the like. A graph may include one or more graph vertices, interconnected by one or more graph edges." *Id.* at 11:51-55. The network graph includes vertices and graph edges where "each graph vertex

may correspond with a network entity included in the network, cloud, environment, or the like, and each graph edge may correspond with a connection between such entities.” *Id.* 11:57-60.

The vertices in the network graph may represent “generic entities, such as are described with respect to S220, imputed generic entities, such as are described with respect to S230, as well as any combination thereof.” *Id.* at 62-64.

53. The '896 patent generates the network graph using network entities. “Network entities 105, as may be included in a cloud platform 104, are entities, systems, devices, components, applications, objects, and the like, configured to operate within the cloud platform 104 and provide various functionalities therein. Specifically, the network entities 105 may be, as examples without limitation, entities configured to process data, send data, or receive data, as well as entities configured to provide various other functionalities, and any combination thereof. The network entities 105 may be configured to connect with various other network entities 105, various external entities, and the like, as well as any combination thereof, for purposes including, without limitation, sending data, receiving data, monitoring data transmissions, monitoring network status and activity, and the like, as well as any combination thereof.” *Id.* at 5:29-43. In addition to network entities, the '896 patent includes imputed entities in the network graph. “Imputed entities are generic entities similar or identical to those described with respect to S220, above, which may be constructed to provide representation of network entities which are integrated into host platforms, or network entities which are shielded from, or not otherwise exposed to, a system configured to execute network analysis processes and methods, including the method described with respect to FIG. 2, where such a system may be, without limitation, the graph analysis system 150 of FIG. 1A.” *Id.* at 10:58-67.

54. The network entities may be collected from a plurality of cloud computing platforms. “A cloud platform 104 may be a commercially-available cloud system, provided on a service basis, such as, as examples and without limitation, Amazon AWS®, Microsoft Azure®, and the like. A cloud platform 104 may be a private cloud, a public cloud, a hybrid cloud, and the like. In addition, a cloud platform 104 may include, without limitation, container orchestration or management systems or platforms such as, as an example and without limitation, a Kubernetes deployment, and the like, as well as any combination thereof.” *Id.* 5:1-11 The ’896 patent discloses using the network graph for “network analysis, traffic analysis, entity querying, graph generation and the like, as well as any combination thereof.” *Id.* 6:48-50. “The graph analysis system 150 may be configured as a physical system, device, or component, as a virtual system, device, or component, or in a hybrid physical-virtual configuration.” *Id.* at 6:55-58. Finally, “the network graph is stored in a graph database.” *Id.* 13:41-42.

55. Orca has infringed and continues to directly infringe one or more claims of the ’896 patent by making, using, selling, offering for sale, and/or importing into the United States without authority or license, the Orca Platform with Attack Path Analysis in violation of 35 U.S.C. § 271(a). Orca’s infringement includes infringement of, for example, claim 1 of the ’896 patent.

56. Claim of the ’896 patent recites:

1. A method for generation of unified graph models for network entities, comprising:
 - collecting, for each network entity of a plurality of network entities, network entity data, wherein the network entity data collected for a network entity includes at least a network entity property, wherein the plurality of network entities are deployed in a plurality of cloud computing platforms;
 - genericizing each of the network entities based on the respective collected network entity data to generate a plurality of generic network entities, wherein a generic network entity includes a generic representation of respective network entities from different cloud computing platforms of the plurality of cloud computing platforms;

generating at least a network graph, wherein the generated network graph is a multi-dimensional data structure providing a representation of the plurality of generic network entities and relations between the generic network entities of the plurality of network entities; and
creating at least one imputed entity, wherein the at least one imputed entity is a generic network entity representing an executed platform functionality, and wherein the executed platform functionality is different than a network entity; and
storing the generated network graph.

57. On information and belief, Orca practices each and every limitation of claim 1 of the '896 patent by and through the use of the Attack Path Analysis.

58. The preamble of claim 1 recites “A method for generation of unified graph models for network entities, comprising:” To the extent the preamble is limiting, Orca practices this step by, for example but not limited to, using its Attack Path Analysis product to generate a unified graph model for network entities. *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/> (“For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and

internal cloud connectivity.”).

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization’s most valuable assets, or ‘crown jewels’, and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the core capabilities of a Cloud-Native Application Protection Platform (CNAPP):

59. Claim 1 further recites “collecting, for each network entity of a plurality of network entities, network entity data, wherein the network entity data collected for a network entity includes at least a network entity property, wherein the plurality of network entities are deployed in a plurality of cloud computing platforms;” Orca’s public blog posts confirm that Orca practices this step by, for example but not limited to, Orca’s Attack Path Analysis “collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.” *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/> (“To fully understand where your organization’s most critical weaknesses are, it is important to view risks as an interrelated chain, rather than just siloed risks. By understanding which combinations are a direct path to your

critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard.”).

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

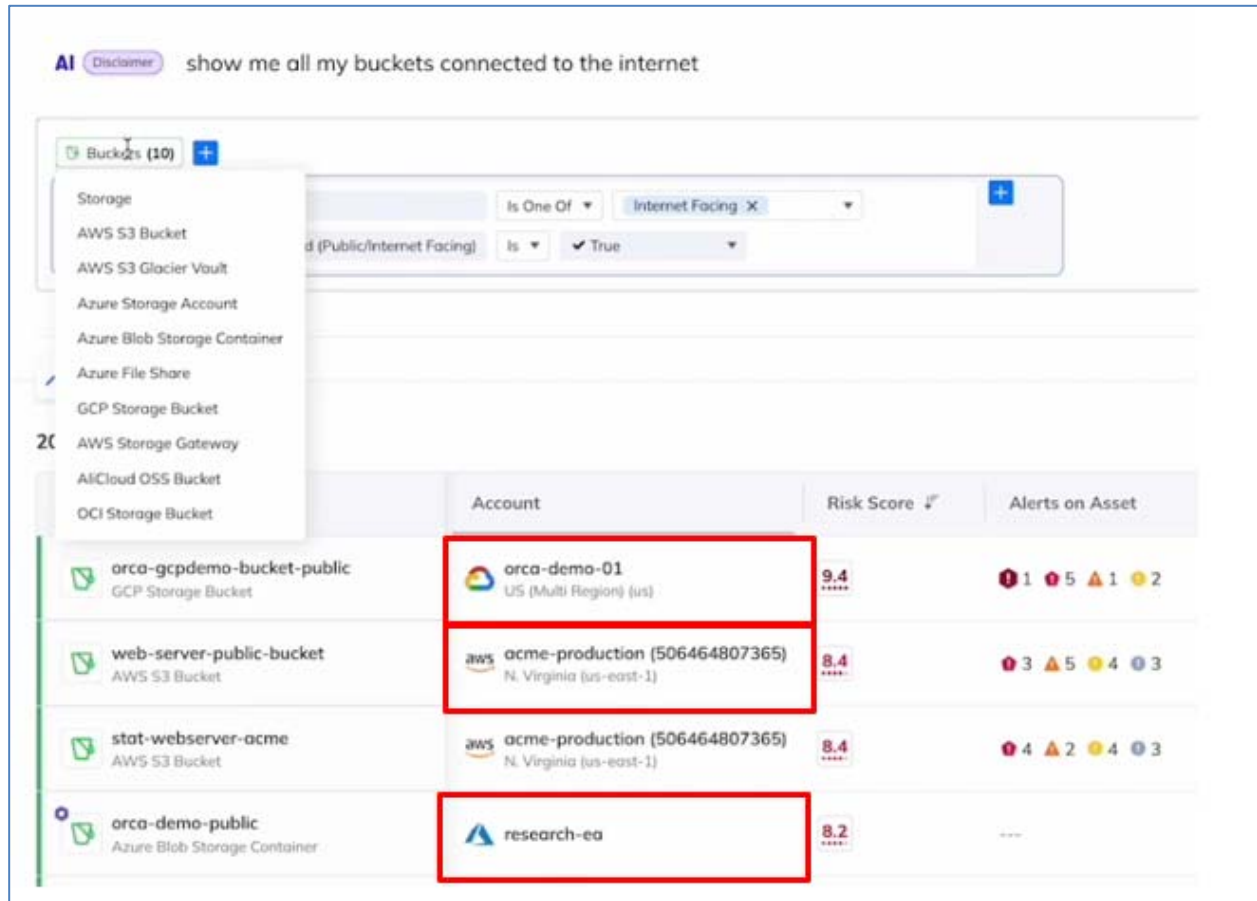
Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the core capabilities of a Cloud-Native Application Protection Platform (CNAPP):

60. Orca collects network entity data across a plurality of cloud computing platforms. See, e.g., *Key Security Capabilities in Kubernetes*, <https://orca.security/resources/blog/kubernetes-security-capabilities-policies/>; *AI-Driven Cloud Security*, <https://orca.security/platform/ai-cloud-security/> (“I’m going to use AI and I’m going to say show me all my buckets connected to the internet. Now we’re taking that query, the natural language that I used, so think about me being someone who perhaps doesn’t know all the different types of buckets that exist across all the different cloud service providers. Well, now I don’t need to know that I’m looking for a GCP storage bucket or an S3 bucket or any other parts

of the storage.”).



61. Claim 1 further recites “genericizing each of the network entities based on the respective collected network entity data to generate a plurality of generic network entities, wherein a generic network entity includes a generic representation of respective network entities from different cloud computing platforms of the plurality of cloud computing platforms;” On information and belief, Orca practices this step by, for example but not limited to, Orca’s Attack Path Analysis “collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.” See, e.g., *Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/> (“To fully understand where your organization’s most critical weaknesses are, it is important to view risks as an interrelated chain,

rather than just siloed risks. By understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard.”). *See Key Security Capabilities in Kubernetes*, <https://orca.security/resources/blog/kubernetes-security-capabilities-policies/>.

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

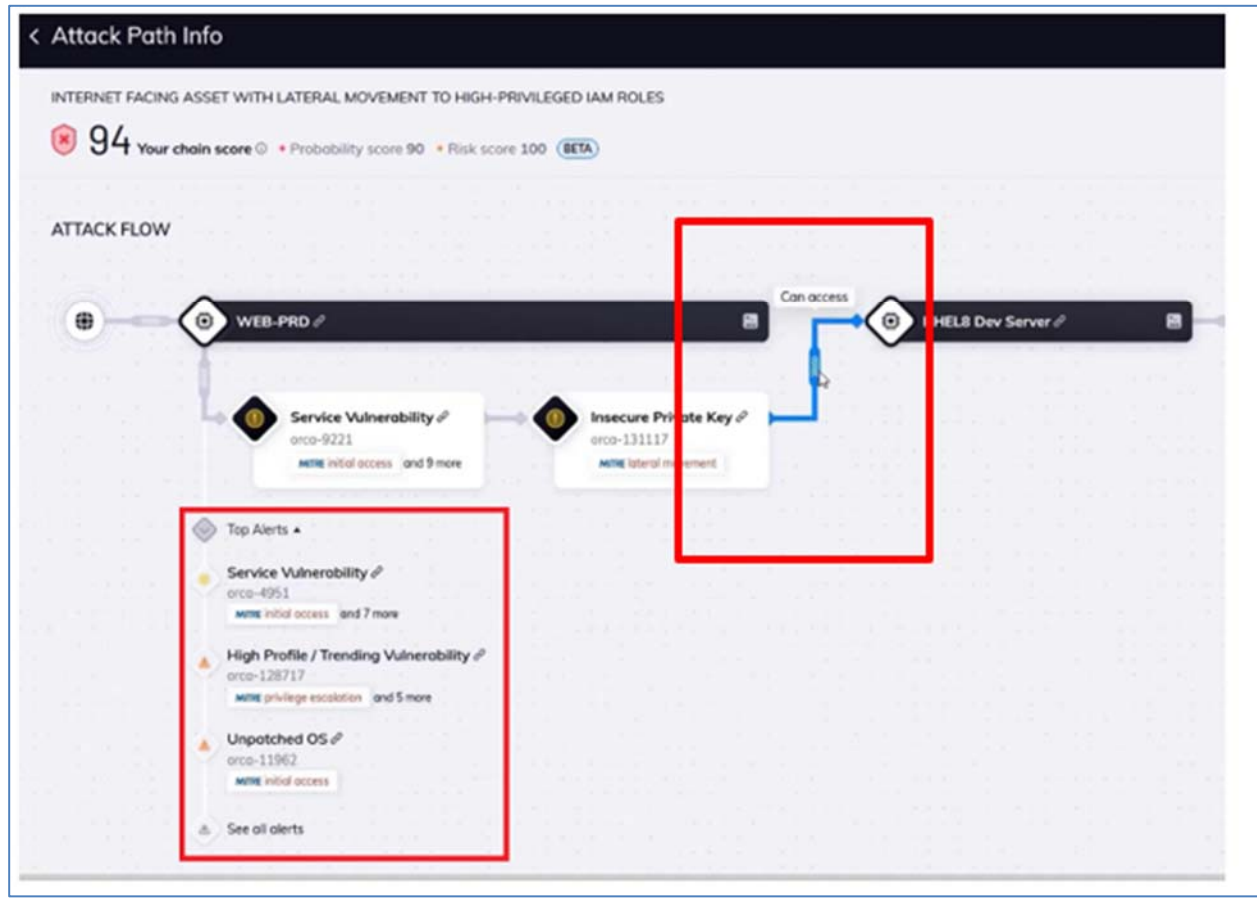
Orca then combines this information with the location of the organization’s most valuable assets, or ‘crown jewels’, and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the core capabilities of a Cloud-Native Application Protection Platform (CNAPP):

As a further example, Orca displays the generic entities in their attack path visualization. *See id.* (“Attack Path visualization of how an attacker can access sensitive company data”); *see, e.g.*,

Orca Bytes: Attack Path Analysis, <https://www.youtube.com/watch?v=MJkO8UfQa-8>.



62. Claim 1 further recites “generating at least a network graph, wherein the generated network graph is a multi-dimensional data structure providing a representation of the plurality of generic network entities and relations between the generic network entities of the plurality of network entities; and. . . .” On information and belief, Orca practices this step as shown by, for example but not limited to, describing operation of Orca’s Attack Path Analysis as “representing attack paths in a visual graph with contextual data on all relevant cloud entities” and by using their “attack path visualization” tool. *See, e.g., Cloud Attack Path Analysis: Work*

Smarter Not Harder, <https://orca.security/resources/blog/cloud-attack-path-analysis/>.

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

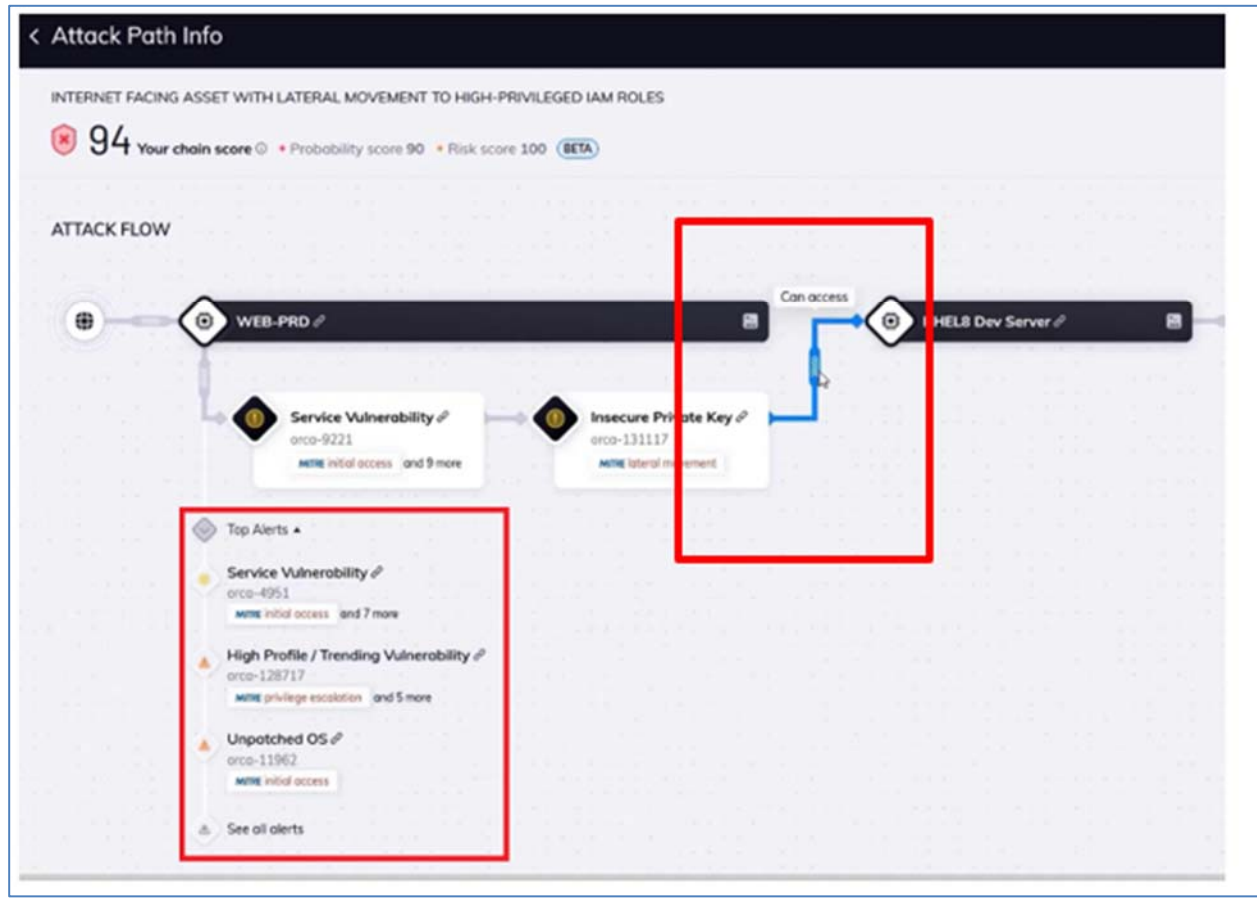
Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates [contextual data](#) on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the [core capabilities of a Cloud-Native Application Protection Platform \(CNAPP\)](#):

63. As a further example, on information and belief, Orca's network graph is a multi-dimensional data structure representing network entities and each entity contains properties,

including but not limited to, entity names and associated alerts.

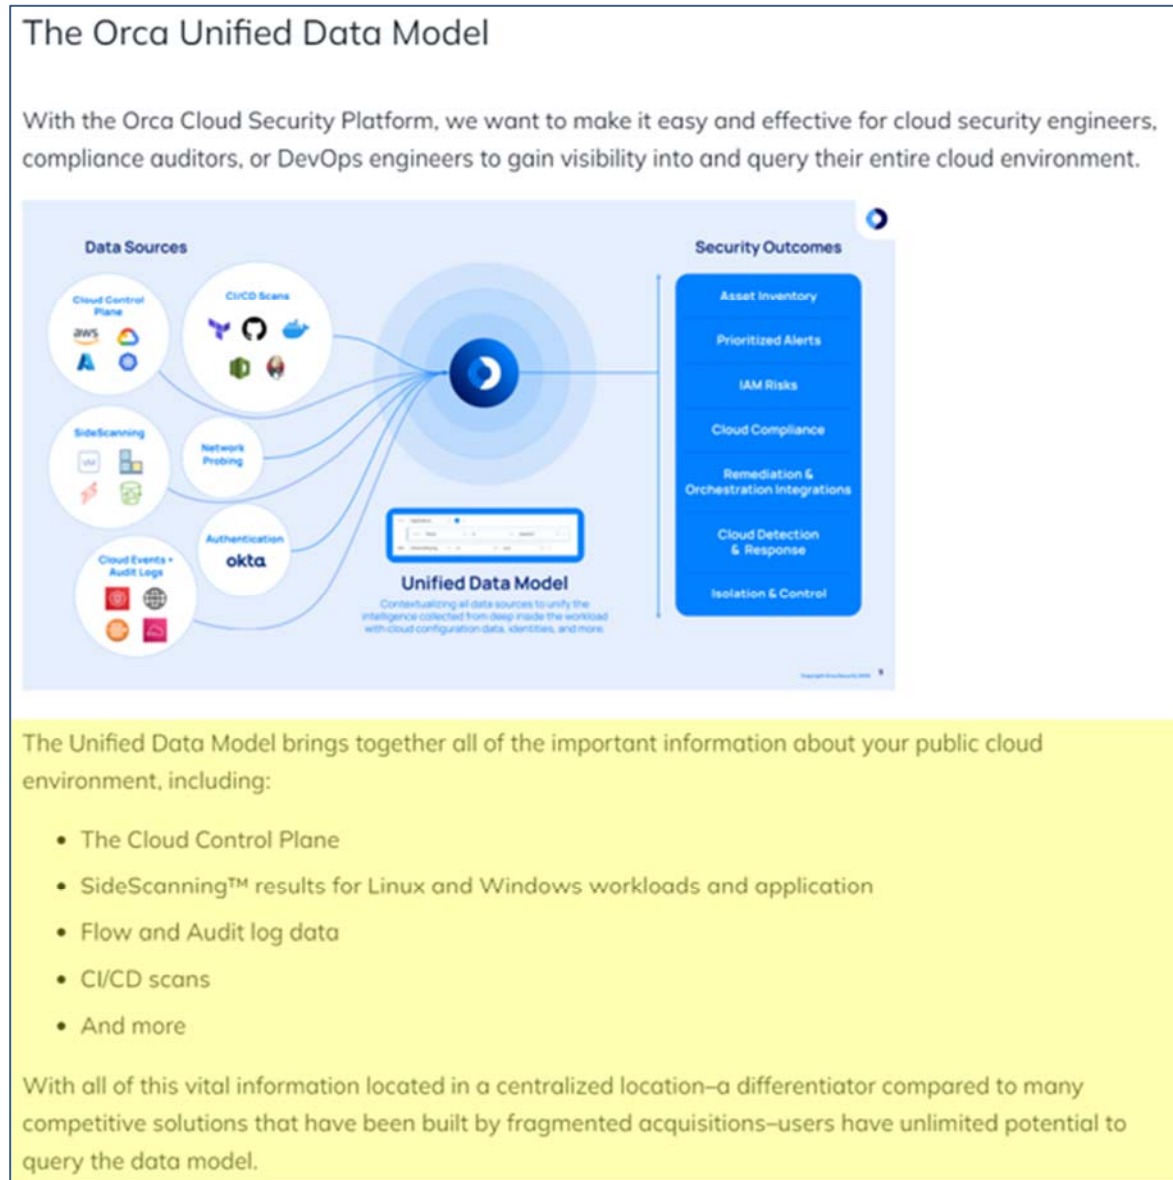


64. Claim 1 further recites “creating at least one imputed entity, wherein the at least one imputed entity is a generic network entity representing an executed platform functionality, and wherein the executed platform functionality is different than a network entity; and” On information and belief, Orca practices this step by, for example but not limited to, including firewall rules in their graph visualization of cloud assets. *See, e.g., Navigating Your Cloud Estate to Understand External Exposure*, <https://orca.security/resources/blog/understanding-external-exposure-with-graph-visualization/>.



65. Finally, claim 1 further recites “storing the generated network graph.” Orca’s public blog posts confirm that Orca practices this step by, for example but not limited to, using Orca’s Unified Data Model to store a network graph. *See, e.g., Cloud Security Simplified: Easily Query Your Entire Cloud Environment*, <https://orca.security/resources/blog/orca-sonar-data-query-builder/> (“The Unified Data Model brings together all of the important information about your public cloud environment . . . ”); and its “Data Security and Posture Management.” *See e.g., Data Security and Posture Management*, <https://orca.security/platform/data-security-and-posture-management-dspm/> (“The Orca Cloud Security Platform performs continuous discovery of data stores across your cloud estate, and alerts to security and compliance risks, without requiring any additional tools. Instead of focusing solely on data security, Orca delivers a comprehensive, context-driven picture of sensitive data exposure, enabling organizations to prioritize risks effectively, reduce alert fatigue, and stay focused on what matters most—from a

single platform.”).



66. On information and belief, Orca was aware of the '896 patent and Orca's infringement thereof prior to these counterclaims at least due to Orca's monitoring of Wiz patents as shown by, in its original complaint in this action, Orca cited and quoted from Wiz's U.S. Patent No. 11,374,982. *See, e.g.*, D.I. 1, ¶ 22. Further, as demonstrated above Orca has repeatedly shown a culture of copying Wiz. This is just one more example of Orca seeing Wiz's success and copying instead of innovating. Moreover, Orca is aware of the '896 patent and

Orca's infringement thereof at least as of the filing of these counterclaims. Accordingly, Orca has and continues to willfully infringe the '896 patent.

67. Orca has induced and continues to induce infringement of one or more claims of the '896 patent by, for example but not limited to, encouraging customers to its Attack Path Analysis in a manner that directly infringes those claims. Despite its knowledge of the existence of the '896 patent, since at least the filing of this Counterclaim, Orca, upon information and belief, continues to encourage, instruct, enable and otherwise cause its customers to use its Attack Path Analysis in a manner that infringes one or more claims of the '896 patent. Upon information and belief, Orca specifically intends that its customers use its Attack Path Analysis in a manner that infringes one or more claims of the '896 patent by, at a minimum, providing instructions and/or support documentation directing customers on how to use its Attack Path Analysis in an infringing manner, in violation of 35 U.S.C. § 271(b). For example, Orca's public blog posts cited above provide instructions and encourage customers to practice all steps of the claimed method stating "To fully understand where your organization's most critical weaknesses are, it is important to view risks as an interrelated chain, rather than just siloed risks. By understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard." *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/>. Further, Orca provides video explanation of Attack Path Analysis stating, "with Orca's new attack path analysis and prioritization capabilities, security teams can now laser focus on a small number of prioritized attack paths or alert on combinations that endanger the company's most critical assets, and every path and link in the path is scored so you can

pinpoint exactly which risks need to be remediated.” *See Orca Bytes: Attack Path Analysis* (Mar. 31, 2022), <https://www.youtube.com/watch?v=MJkO8UfQa-8>.

68. Orca has contributed and continues to contribute to the infringement of one or more claims of the ’896 patent. Upon information and belief, Orca knows that the Accused Product is especially made and/or adapted for users to infringe one or more claims of the ’896 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use. Because Orca included features, such as for example but not limited to Attack Path Analysis, in Orca’s products, Orca intends for customers to use it. Upon information and belief, the Attack Path Analysis feature has no suitable use that is non-infringing, and therefore Orca intends for customers to use its Attack Path Analysis in an infringing manner. Orca’s sales of products including Attack Path Analysis constitute contributory infringement in violation of 35 U.S.C. § 271(c).

COUNTERCLAIM III
(INFRINGEMENT OF U.S. PATENT NO. 11,936,693)

69. Wiz is the sole and exclusive owner, by assignment, of all rights, title and interest in U.S. Patent No. 11,936,693 (the “’693 patent”), entitled “System and Method for Applying a Policy on a Network Path.” The ’693 patent was duly and legally issued by the U.S. Patent and Trademark Office on Mar. 19, 2024. The named inventors of the ’693 patent are Roy Reznik, Matilda Lidgi, Shai Keren, and Eliran Marom. A copy of the ’693 patent is attached as Exhibit C.

70. The ’693 patent generally relates to applying a policy on a network path to a reachable resource in a cloud computing environment. *See* ’693 patent at 2:44-54. The policy may include a conditional rule which, if not met, initiates a mitigating action. *Id.* at 2:54-56.

71. The '693 patent discloses “initiating active inspection for each network path of a plurality of network paths; storing an indicator to indicate that a first network path of the plurality of network paths is a valid path, in response to determining that the reachable resource is accessible from the external network; and applying the policy on the first network path.” *Id.* at 2:64-3:3. The '693 patent further discloses “initiating the mitigation action on the reachable resource . . . where the mitigation action includes any one of: revoking access to the reachable resource, revoking access from the reachable resource, closing a port of the reachable resource, generating a notification, generating an alert, and any combination thereof.” *Id.* at 3:11-16.

72. Orca has infringed and continues to directly infringe one or more claims of the '693 patent by making, using, selling, offering for sale, and/or importing into the United States without authority or license, the Orca Platform with Attack Path Analysis and Auto Remediation in violation of 35 U.S.C. § 271(a). Orca's infringement includes infringement of, for example, claim 1 of the '693 patent.

73. Claim 1 of the '693 patent recites:

1. A method for applying a policy on a network path, comprising:
 - selecting a reachable resource having a network path to access the reachable resource, wherein the reachable resource is a cloud object deployed in a cloud computing environment, having access to an external network which is external to the cloud computing environment;
 - actively inspecting the network path to determine if the network path of the reachable resource is accessible from the external network;
 - storing an indicator to indicate that the network path is a valid path, in response to determining that the reachable resource is accessible from the external network;
 - applying a policy on the valid path, wherein the policy includes a conditional rule;
 - initiating a mitigation action, in response to determining that the conditional rule is not met; and
 - applying the policy on another network path, in response to determining that the conditional rule is met.

74. On information and belief, Orca practices each and every limitation of claim 1 of the '693 patent by and through the use of the Attack Path Analysis and Auto-Remediation.

75. The preamble of claim 1 recites “[a] method for applying a policy on a network path, comprising:” To the extent the preamble is limiting, Orca practices this step by, for example but not limited to, using its Attack Path Analysis product to tag network paths. *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/> (“For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.”); *see also id* (“Each attack vector in the path includes tags for easy identification and filtering, including applicable MITRE ATT&CK categories.”).

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the core capabilities of a Cloud-Native Application Protection Platform (CNAPP):

Each attack vector in the path includes tags for easy identification and filtering, including applicable MITRE ATT&CK categories.



Orca shows detailed information for each step

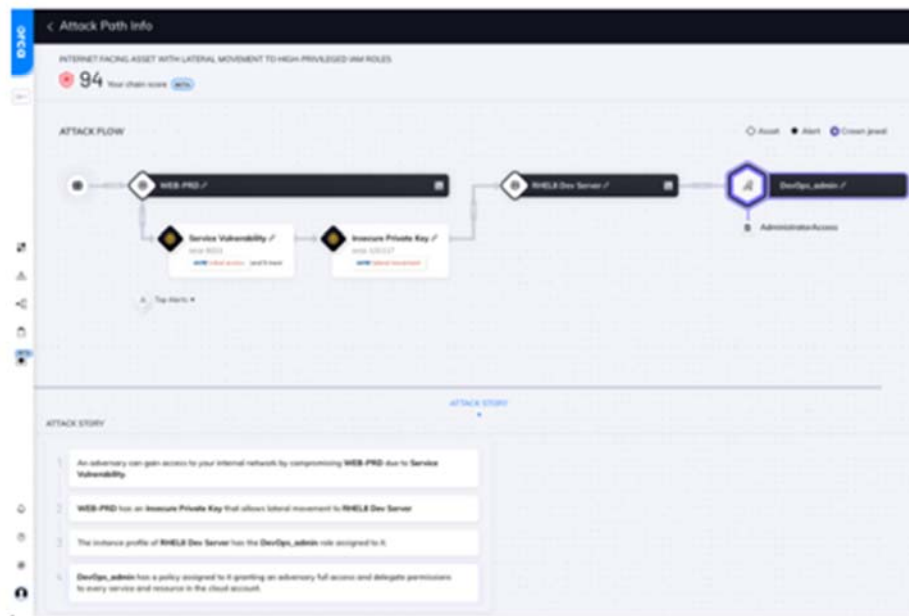
76. Claim 1 further recites “selecting a reachable resource having a network path to access the reachable resource, wherein the reachable resource is a cloud object deployed in a

cloud computing environment, having access to an external network which is external to the cloud computing environment.” Orca’s public blog posts confirm that Orca practices this step by, for example but not limited to, using Orca’s Attack Path Analysis. Orca scans and selects resources that are external or internet facing in a cloud environment. *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis> (“For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.”).

77. Orca further boasts that “[b]y understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard.” *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*,

<https://orca.security/resources/blog/cloud-attack-path-analysis/>.

To fully understand where your organization's most critical weaknesses are, it is important to view risks as an interrelated chain, rather than just siloed risks. By understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard.



Orca shows the steps an attacker can take to reach the company's crown jewels

78. Further, Orca considers “Accessibility” as a factor when providing Attack Path Scoring and Prioritization. Accessibility includes whether or not a resource is “internet facing.” See, e.g., *id.* (“Accessibility: Is the attack path Internet facing? How easy is it to exploit the

initial entry point to the attack path?”).

Orca Security Attack Path Scoring and Prioritization

Orca assigns an overall score (from 0 to 99) to each attack path and scores each attack vector that makes up the attack path. To calculate the score, Orca uses an advanced algorithm that takes the following factors into account:

1. **Severity:** How severe is the underlying security issue? For example, what type of threat is it, how likely is it to be exploited, and what is the CVSS score?
2. **Accessibility:** Is the attack path Internet facing? How easy is it to exploit the initial entry point to the attack path?
3. **Lateral Movement Risk:** Is there [lateral movement risk](#)? If so, how easy is it to exploit, and how many hops do you need to reach your end goal?
4. **Access Level:** Does the risk provide read only access, or read and write access? If the risk allows privilege exploitation, what level of privileges does it allow?
5. **Business Impact:** How critical is the target that the attack path exposes? Is it Personal Identifiable Information (PII), or other sensitive information such as intellectual property? Is it a production server that is essential to the business?



How Orca Security scores attack paths

79. Claim 1 further recites “actively inspecting the network path to determine if the network path of the reachable resource is accessible from the external network;” Orca’s public blog posts confirm that Orca practices this step by, for example but not limited to, using Orca’s Attack Path Analysis. Orca states that “it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including

information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.” Further this model is used to determine the accessibility or reachability of an object. *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/> (“Orca assigns an overall score (from 0 to 99) to each attack path and scores each attack vector that makes up the attack path. To calculate the score, Orca uses an advanced algorithm that takes the following factors into account: . . . 2. Accessibility: Is the attack path Internet facing? How easy is it to exploit the initial entry point to the attack path?”).

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the core capabilities of a Cloud-Native Application Protection Platform (CNAPP):

Orca Security Attack Path Scoring and Prioritization

Orca assigns an overall score (from 0 to 99) to each attack path and scores each attack vector that makes up the attack path. To calculate the score, Orca uses an advanced algorithm that takes the following factors into account:

1. **Severity:** How severe is the underlying security issue? For example, what type of threat is it, how likely is it to be exploited, and what is the CVSS score?
2. **Accessibility:** Is the attack path Internet facing? How easy is it to exploit the initial entry point to the attack path?
3. **Lateral Movement Risk:** Is there [lateral movement risk](#)? If so, how easy is it to exploit, and how many hops do you need to reach your end goal?
4. **Access Level:** Does the risk provide read only access, or read and write access? If the risk allows privilege exploitation, what level of privileges does it allow?
5. **Business Impact:** How critical is the target that the attack path exposes? Is it Personal Identifiable Information (PII), or other sensitive information such as intellectual property? Is it a production server that is essential to the business?



How Orca Security scores attack paths

80. Claim 1 further recites “storing an indicator to indicate that the network path is a valid path, in response to determining that the reachable resource is accessible from the external network;” Orca’s public blog posts confirm that Orca practices this step. For example, Orca states that “it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.”

Further this model is used to determine the accessibility or reachability of an object. *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder,*

<https://orca.security/resources/blog/cloud-attack-path-analysis/> (“Orca assigns an overall score (from 0 to 99) to each attack path and scores each attack vector that makes up the attack path. To calculate the score, Orca uses an advanced algorithm that takes the following factors into account: . . . 2. Accessibility: Is the attack path Internet facing? How easy is it to exploit the initial entry point to the attack path? . . . Each attack vector in the path includes tags for easy identification and filtering, including applicable MITRE ATT&CK categories”).

Cloud Attack Path Analysis Automatically Correlates Multiple Alerts into One Interactive View

Orca defines Attack Path Analysis as the automatic identification of risk combinations that create dangerous attack paths that can be exploited by attackers. This includes representing attack paths in a visual graph with contextual data on all relevant cloud entities and their risks across vulnerability status, misconfiguration risks, trust and authorization, and data as well as the relations between them.

Orca then combines this information with the location of the organization's most valuable assets, or 'crown jewels', and assigns each attack path a Business Impact Score. This allows security teams to immediately understand which attack paths are the most critical to the business, so they can remediate those first.

For attack path analysis to be truly beneficial, it is essential that the cloud security platform utilizes a unified data model that collects and correlates contextual data on each asset, including information on potential risks in the cloud workload and configuration as well as external and internal cloud connectivity.

Industry analysts are increasingly recognizing the value and need for more detailed attack path analysis as part of cloud security solutions. Gartner lists attack path analysis as one of the core capabilities of a Cloud-Native Application Protection Platform (CNAPP):

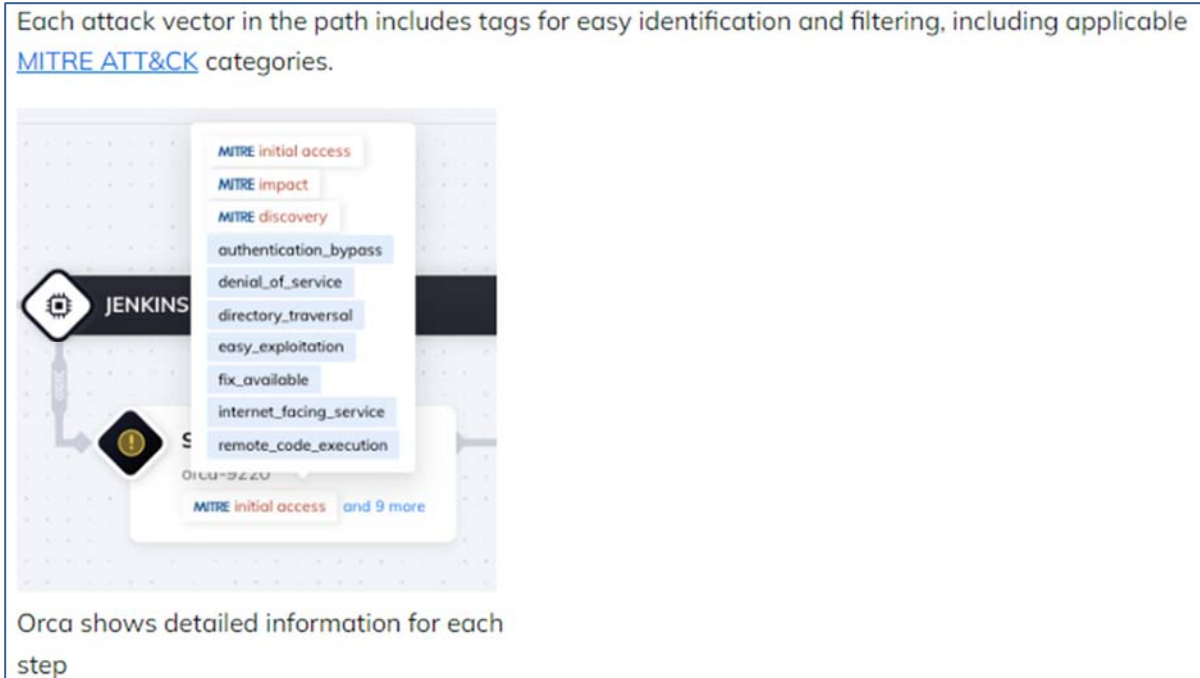
Orca Security Attack Path Scoring and Prioritization

Orca assigns an overall score (from 0 to 99) to each attack path and scores each attack vector that makes up the attack path. To calculate the score, Orca uses an advanced algorithm that takes the following factors into account:

1. **Severity:** How severe is the underlying security issue? For example, what type of threat is it, how likely is it to be exploited, and what is the CVSS score?
2. **Accessibility:** Is the attack path Internet facing? How easy is it to exploit the initial entry point to the attack path?
3. **Lateral Movement Risk:** Is there [lateral movement risk](#)? If so, how easy is it to exploit, and how many hops do you need to reach your end goal?
4. **Access Level:** Does the risk provide read only access, or read and write access? If the risk allows privilege exploitation, what level of privileges does it allow?
5. **Business Impact:** How critical is the target that the attack path exposes? Is it Personal Identifiable Information (PII), or other sensitive information such as intellectual property? Is it a production server that is essential to the business?



How Orca Security scores attack paths



See also, *Data Security and Posture Management*, <https://orca.security/platform/data-security-and-posture-management-dspm/> (“Orca’s DSPM dashboard provides data security teams with an overview of their cloud data stores, sensitive data, and security and compliance alerts. Orca scans managed, unmanaged, and shadow data, giving security teams wide and deep visibility

into where their data resides.”).

Discover and classify data in your cloud

Orca's DSPM dashboard provides data security teams with an overview of their cloud data stores, [sensitive data](#), and security and compliance alerts. Orca scans managed, unmanaged, and shadow data, giving security teams wide and deep visibility into where their data resides.

- ✓ Get a multi-cloud inventory of data storage assets—including databases, and files in virtual machines, storage buckets, and containers.
- ✓ Know which data stores contain sensitive data and of what type—including PII, [PCI](#), PHI, or financial information—for both security and regulatory purposes.
- ✓ Leverage interactive graphs that show the location and relationship between data stores and other cloud entities.

81. Claim 1 further recites “applying a policy on the valid path, wherein the policy includes a conditional rule. . . .” On information and belief, Orca practices this step through the use of Auto Remediation. Orca states for example, that it applies “auto remediation” policies on assets with common and complex security alerts. *See, e.g. Manage Cloud Security Risks with Auto-Remediation*, <https://orca.security/resources/blog/manage-security-risks-auto-remediation/>.

(“Orca Security introduces Automatic Remediation: a way to quickly resolve common and complex security alerts, such as an Unencrypted S3 Bucket or Security group with permissive access, reducing friction between different groups in the organization and increasing productivity. With Orca Automatic Remediation, you can configure automation rules which remediate alerts as they are detected or click the “Auto-Remediation” button on a specific alert (indicated by a green magic stick icon).”).

Automatic Remediation as an Approach to Better, Faster Security

One of the most complex parts of the alert life cycle is the manual remediation process; usually, it will include either step-by-step mitigation instructions on the console or copying and pasting command after command to the CLI.

However, an excessive number of alerts can quickly become unmanageable. To assist with this challenge, Orca Security introduces Automatic Remediation: a way to quickly resolve common and complex security alerts, such as an Unencrypted S3 Bucket or Security group with permissive access, reducing friction between different groups in the organization and increasing productivity.

With Orca Automatic Remediation, you can configure automation rules which remediate alerts as they are detected or click the “Auto-Remediation” button on a specific alert (indicated by a green magic stick icon).



See also, <https://orca.security/resources/video/auto-remediation-demo/> (Demonstrating setting up

a policy on an internet facing asset.)

The screenshot displays the Orca security dashboard for an asset named 'acme-dev2948', which is an AWS S3 Bucket. The dashboard is divided into several sections:

- ALERTS ON ASSET:** Shows a summary of alerts: 1 Compromise (red), 0 Imminent Compromise (orange), and 3 Hazardous (yellow). It also indicates 2 Data at risk, 1 Data protection, and 1 other alert.
- ASSET INFORMATION:**
 - Alerts on asset:** A list of top alerts. The first alert is 'Malware' from 'jeicar.com' seen 8 hours ago. The second and third alerts are 'S3 Bucket Should Enforce H...' and 'S3 Bucket Allows Public REA...' respectively, both seen 8 hours ago. The fourth alert is 'Bucket Allows Unencrypted ...' also seen 8 hours ago. A tooltip 'Auto remediation available' is visible over the third alert.
 - Attack vector:** A diagram showing a path from 'public' to 'acme-dev2948'.
- Identity:**
 - Alerts on asset:** 1 red, 3 yellow, 9 blue.
 - Asset type:** AwsS3Bucket.
 - Public:** true.
 - Internet facing:** Yes.
 - Cloud account:** acme-production...
 - Attack paths:** 6 Low Impact.
 - Cloud account ID:** 506464807365.
 - Status:** Enabled.
- Additional Information:**
 - Create Time:** 2 years ago (2020, Oct 16).
 - Up Time:** 6 hours ago (2022, Sep 28, 03:35).

82. Claim 1 further recites “initiating a mitigation action, in response to determining that the conditional rule is not met; and” Orca’s public blog posts confirm that Orca practices this step. Orca states for example, that it applies “auto remediation” policies on assets with common and complex security alerts when security rules are not met. These remediation actions can include hardening permissive access rules to assets and blocking specific ports. *See, e.g. Manage Cloud Security Risks with Auto-Remediation,*

<https://orca.security/resources/blog/manage-security-risks-auto-remediation/> (“Orca Security introduces Automatic Remediation: a way to quickly resolve common and complex security alerts, such as an Unencrypted S3 Bucket or Security group with permissive access, reducing friction between different groups in the organization and increasing productivity. With Orca Automatic Remediation, you can configure automation rules which remediate alerts as they are detected or click the “Auto-Remediation” button on a specific alert (indicated by a green magic stick icon) For example, you can automatically harden permissive access on insecure security group rules and block specific ports while creating a Jira ticket notifying your DevOps team with more details. Our remediation capabilities give you the option to select the action based on your requirements.”).

Automatic Remediation as an Approach to Better, Faster Security

One of the most complex parts of the alert life cycle is the manual remediation process; usually, it will include either step-by-step mitigation instructions on the console or copying and pasting command after command to the CLI.

However, an excessive number of alerts can quickly become unmanageable. To assist with this challenge, Orca Security introduces Automatic Remediation: a way to quickly resolve common and complex security alerts, such as an Unencrypted S3 Bucket or Security group with permissive access, reducing friction between different groups in the organization and increasing productivity.

With Orca Automatic Remediation, you can configure automation rules which remediate alerts as they are detected or click the “Auto-Remediation” button on a specific alert (indicated by a green magic stick icon).

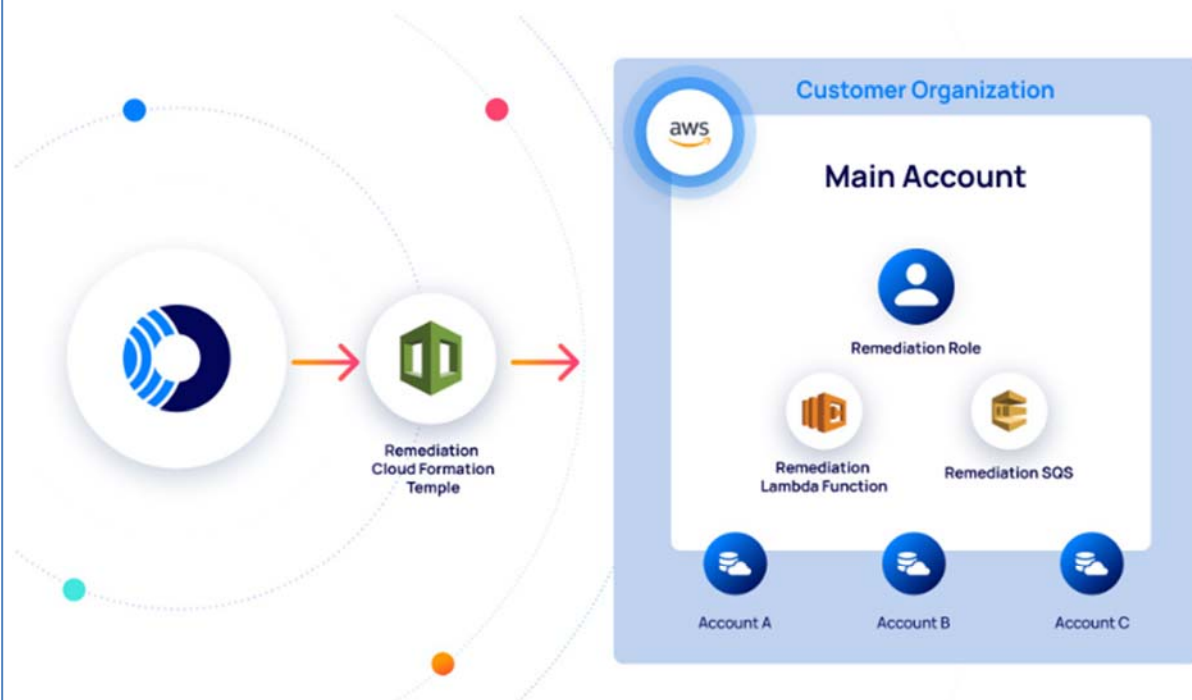


How Does Orca Auto-Remediation Work?

Orca's Auto-Remediation is yet another AWS CloudFormation stack deployed in your environment, which means that you do not need to provide write access to Orca.


Instead, the remediation solution infrastructure resides in the customer's environment.

The Orca Cloud Security Platform sends remediation instructions to an AWS SQS Queue, triggering a Lambda function. The function then calls the appropriate action to remediate the alert(s), as shown in the diagram below.





Using Orca's Auto-Remediation, you can quickly and easily remediate security issues in your cloud environment. For example, you can automatically harden permissive access on insecure security group rules and block specific ports while creating a Jira ticket notifying your DevOps team with more details. Our remediation capabilities give you the option to select the action based on your requirements.

See also, <https://orca.security/resources/video/auto-remediation-demo/> (Demonstrating applying auto remediation to a S3 bucket that allows public read access).



Data at risk

S3 Bucket Allows Public READ_...



Take action

ALERT INFO

ASSET INFO

Status

Open

Integration


Summary

Ensure that your S3 buckets content permissions details cannot be viewed by anonymous users in order to protect against unauthorized access. An S3 bucket that grants READ_ACP (view permissions) access to everyone can

SHOW MORE

Remediation

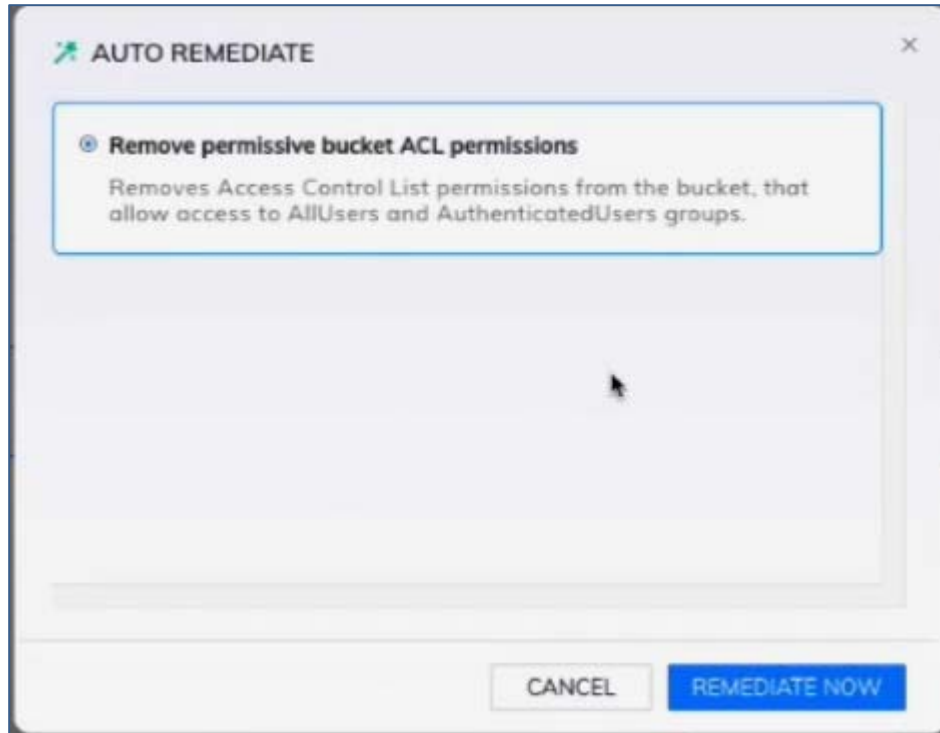
Auto

 Choose auto remediation

or

Manual

Change the acme-dev2948 bucket policy to block public READ_ACP access



83. Finally claim 1 recites “applying the policy on another network path, in response to determining that the conditional rule is met.” On information and belief, Orca practices this step. Orca states for example, that it “Automate[s] routine remediation tasks and processes,” “Reduce[s] redundancy,” and that “you can configure automation rules which remediate alerts as they are detected.” *See, e.g. Manage Cloud Security Risks with Auto-Remediation*, <https://orca.security/resources/blog/manage-security-risks-auto-remediation/>.

84.

85. applies “auto remediation” policies on assets with common and complex security alerts. These remediation actions can include hardening permissive access rules to assets and blocking specific ports in “security groups” impacting multiple network paths. Orca specifically states that applying these policies and mitigation actions across a multi-cloud environment through automation “is the way to improve accuracy, reduce redundancy, and reduce cost and validation time.” This demonstrates how Orca applies remediation polices across network paths.

See, e.g. Manage Cloud Security Risks with Auto-Remediation,

<https://orca.security/resources/blog/manage-security-risks-auto-remediation/>; *see also id.* (“Orca Security introduces Automatic Remediation: a way to quickly resolve common and complex security alerts, such as an Unencrypted S3 Bucket or Security group with permissive access, reducing friction between different groups in the organization and increasing productivity. With Orca Automatic Remediation, you can configure automation rules which remediate alerts as they are detected or click the “Auto-Remediation” button on a specific alert (indicated by a green magic stick icon) . . . For example, you can automatically harden permissive access on insecure security group rules and block specific ports while creating a Jira ticket notifying your DevOps team with more details. Our remediation capabilities give you the option to select the action based on your requirements . . . Automation is the way to go to improve accuracy, reduce redundancy, and reduce cost and validation time.”).

Automatic Remediation as an Approach to Better, Faster Security

One of the most complex parts of the alert life cycle is the manual remediation process; usually, it will include either step-by-step mitigation instructions on the console or copying and pasting command after command to the CLI.

However, an excessive number of alerts can quickly become unmanageable. To assist with this challenge, Orca Security introduces Automatic Remediation: a way to quickly resolve common and complex security alerts, such as an Unencrypted S3 Bucket or Security group with permissive access, reducing friction between different groups in the organization and increasing productivity.

With Orca Automatic Remediation, you can configure automation rules which remediate alerts as they are detected or click the “Auto-Remediation” button on a specific alert (indicated by a green magic stick icon).

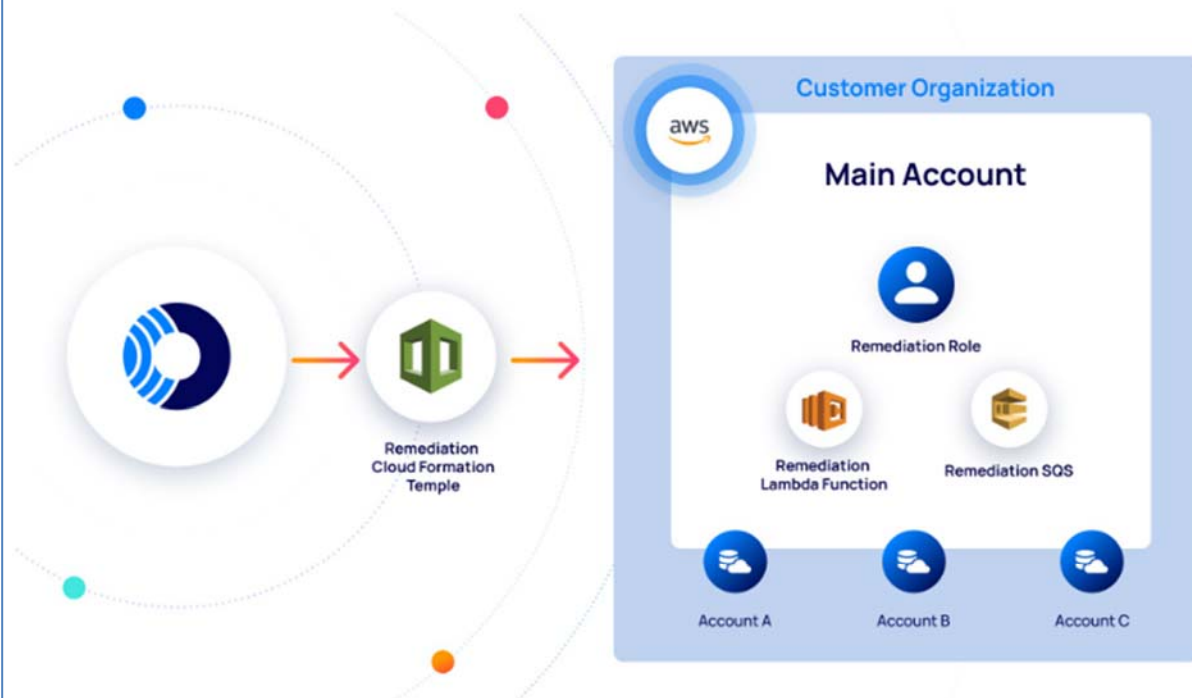


How Does Orca Auto-Remediation Work?

Orca's Auto-Remediation is yet another AWS CloudFormation stack deployed in your environment, which means that you do not need to provide write access to Orca.

Instead, the remediation solution infrastructure resides in the customer's environment.

The Orca Cloud Security Platform sends remediation instructions to an AWS SQS Queue, triggering a Lambda function. The function then calls the appropriate action to remediate the alert(s), as shown in the diagram below.



Using Orca's Auto-Remediation, you can quickly and easily remediate security issues in your cloud environment. For example, you can automatically harden permissive access on insecure security group rules and block specific ports while creating a Jira ticket notifying your DevOps team with more details. Our remediation capabilities give you the option to select the action based on your requirements.

Why Use Orca's Auto-Remediation?

Considering the intricacy of your multi-cloud environments, managing these environments and applications becomes more daunting and complicates your current operational challenges. Automation is the way to go to improve accuracy, reduce redundancy, and reduce cost and validation time.

Auto-Remediation is a self-healing workflow that triggers and responds to alerts or events by executing actions that can prevent or fix the problem. Orca is an event-driven application that uses event-driven automation to resolve policy violations. The Auto-remediation can trigger a serverless function to remediate alerts detected as a result of misconfiguration.

With Orca's Auto-Remediation, your Mean Time to Remediation (MTTR) will be at the bare minimum, thereby improving your security posture and compliance requirements.

Implementing a thorough security framework like [Orca's agentless security solution](#) is the first step in automating your cloud security. Learn how to create custom alerts from queries and integrate these into existing remediation workflows with Orca's platform. Read our [case studies](#) to see how we benefit our customers, or [watch a demo](#) to witness Orca in action. You can also sign up for a [free, no-obligation risk assessment](#) to get started today!

86. Orca is aware of the '693 patent and Orca's infringement thereof at least as of the filing of these counterclaims. Moreover, Orca has a culture of copying Wiz, as explained above. Accordingly, Orca has and continues to willfully infringement the '693 patent.

87. Orca has induced and continues to induce infringement of one or more claims of the '693 patent by encouraging customers to use its Attack Path Analysis and Auto Remediation in a manner that directly infringes those claims. Despite its knowledge of the existence of the '693 patent, since at least the filing of this Counterclaim, Orca, upon information and belief, continues to encourage, instruct, enable and otherwise cause its customers to use its Attack Path Analysis in a manner that infringes one or more claims of the '693 patent. Upon information and belief, Orca specifically intends that its customers use its Attack Path Analysis and Auto Remediation in a manner that infringes one or more claims of the '693 patent by, at a minimum, providing instructions and/or support documentation directing customers on how to use its

Attack Path Analysis and Auto Remediation in an infringing manner, in violation of 35 U.S.C. § 271(b). For example, Orca’s public blog posts cited above provide instructions and encourage customers to practice all steps of the claimed method stating “To fully understand where your organization’s most critical weaknesses are, it is important to view risks as an interrelated chain, rather than just siloed risks. By understanding which combinations are a direct path to your critical assets, security teams can operate strategically by making sure that the risks that break the attack path are remediated first. Orca Security does just that with its new Attack Path Analysis dashboard.” *See, e.g., Cloud Attack Path Analysis: Work Smarter Not Harder*, <https://orca.security/resources/blog/cloud-attack-path-analysis/>. Further, Orca provides instructions on security risk auto remediation stating: “Orca Security introduces Automatic Remediation: a way to quickly resolve common and complex security alerts, such as an Unencrypted S3 Bucket or Security group with permissive access, reducing friction between different groups in the organization and increasing productivity. With Orca Automatic Remediation, you can configure automation rules which remediate alerts as they are detected or click the “Auto-Remediation” button on a specific alert (indicated by a green magic stick icon).” *See, e.g., Manage Cloud Security Risks with Auto-Remediation*, <https://orca.security/resources/blog/manage-security-risks-auto-remediation/>.

88. Orca has contributed and continues to contribute to the infringement of one or more claims of the ’693 patent. Upon information and belief, Orca knows that its Attack Path Analysis and Auto Remediation are especially made and/or adapted for users to infringe one or more claims of the ’693 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use. Because Orca included the features, such as for example but not limited to, Attack Path Analysis and Auto-Remediation, in Orca’s products, Orca intends for

customers to use it. Upon information and belief, Orca's Attack Path Analysis feature has no suitable use that is non-infringing, and therefore Orca intends for customers to use this feature in an infringing manner. Orca's sales of products including its Attack Path Analysis and Auto Remediation constitute contributory infringement in violation of 35 U.S.C. § 271(c).

COUNTERCLAIM IV
(INFRINGEMENT OF U.S. PATENT NO. 12,001,549)

89. Wiz is the sole and exclusive owner, by assignment, of all rights, title and interest in U.S. Patent No. 12,001,549 (the "'549 patent"), entitled "Cybersecurity Incident Response Techniques Utilizing Artificial Intelligence." The '549 patent was duly and legally issued by the U.S. Patent and Trademark Office on June 4, 2024. The named inventors of the '549 patent are Alon Schindel, Barak Sharoni, Amitai Cohen, Ami Luttwak, Roy Reznik, and Yinon Costica. A copy of the '549 patent is attached as Exhibit D.

90. The '549 patent generally relates to providing a cybersecurity incident response to an incident based on a cybersecurity event and generating a prompt for a large language model to generate a query on a security database for a mitigation action. *See* '549 patent at Abstract. The patent provides a method "where the incident input includes any one of: a query, a statement, and a combination thereof." *Id.* at 4:43-44.

91. The '549 patent discloses "providing the received incident input into a large language model (LLM)[.]" *Id.* at 2:35-36. The disclosed LLM is trained on "a data schema utilized in representing the computing environment, incident data classified to a scenario, the plurality of scenarios, and any combination thereof." *Id.* at 2:38-41. The '549 patent further discloses "utilizing the LLM to generate an explanation of a security finding." *Id.* at 2:55-56. From this finding, the system may generate and execute a query in a security database. *Id.* at 4:23-41.

92. Finally, the '549 patent discloses that the “[s]ystem may also initiate a mitigation action based on a result of the executed query.” *Id.* at 3:26-28. “A mitigation action includes generating a notification, generating an alert, updating an alert, generating a severity score, updating a severity score, generating a ticket, generating a risk score, updating a risk score, initiating a remediation action, initiating an incident response, a combination thereof, and the like.” *Id.* at 17:15-21.

93. Orca has infringed and continues to directly infringe one or more claims of the '549 patent by making, using, selling, offering for sale, and/or importing into the United States without authority or license, the Orca Platform with AI-Driven Cloud Security in violation of 35 U.S.C. § 271(a). Orca's infringement includes infringement of, for example, claim 1 of the '549 patent.

94. Claim 1 of the '549 patent recites:

1. A method for providing cybersecurity incident response, comprising:
 receiving an incident input based on a cybersecurity event;
 generating a prompt for a large language model (LLM) based on the received incident input;
 configuring the LLM to generate an output based on the generated prompt;
 mapping the received incident input into a scenario of a plurality of scenarios based on the output of the LLM, wherein each scenario is associated with an incidence response;
 generating a query based on the received incident input and the mapped scenario;
 executing the query on a security database, the security database including a representation of a computing environment; and
 initiating a mitigation action based on a result of the executed query.

95. On information and belief, Orca practices each and every limitation of claim 1 of the '549 patent by and through the use of AI-Driven Cloud Security.

96. The preamble of claim 1 recites “[a] method for providing cybersecurity incident response, comprising:” To the extent the preamble is limiting, Orca practices this step by, for example but not limited to, using its AI-Driven Cloud Security features to provide a

cybersecurity incident response. *See, e.g., AI-Driven Cloud Security,*

<https://orca.security/platform/ai-cloud-security/> (“Simplify investigations and accelerate remediation with built-in generative AI.”).



97. Claim 1 further recites “receiving an incident input based on a cybersecurity event. . . .” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, generating alerts in response to a cyber security incident. *See, e.g., AI-Driven Cloud Security, https://orca.security/platform/ai-cloud-security/* (Orca’s dynamic alert and asset descriptions greatly simplify investigations, summarizing all the important information that teams need to know in an easily consumable manner, reducing investigation time and improving Mean Time To Remediation (MTTR)).

AI-generated alert and asset descriptions

Orca's dynamic alert and asset descriptions greatly simplify investigations, summarizing all the important information that teams need to know in an easily consumable manner, reducing investigation time and improving Mean Time To Remediation (MTTR).

- ✓ For assets, Orca summarizes which risks are found and of what severity, how many attack paths they are part of, whether the asset is Internet-facing, running or paused, and more.
- ✓ For alerts, Orca explains what the risk is, when it was first found, if it is exploitable, whether there's a fix, how an attacker could abuse it, and more.
- ✓ Where applicable, descriptions contain links to other resources with more information

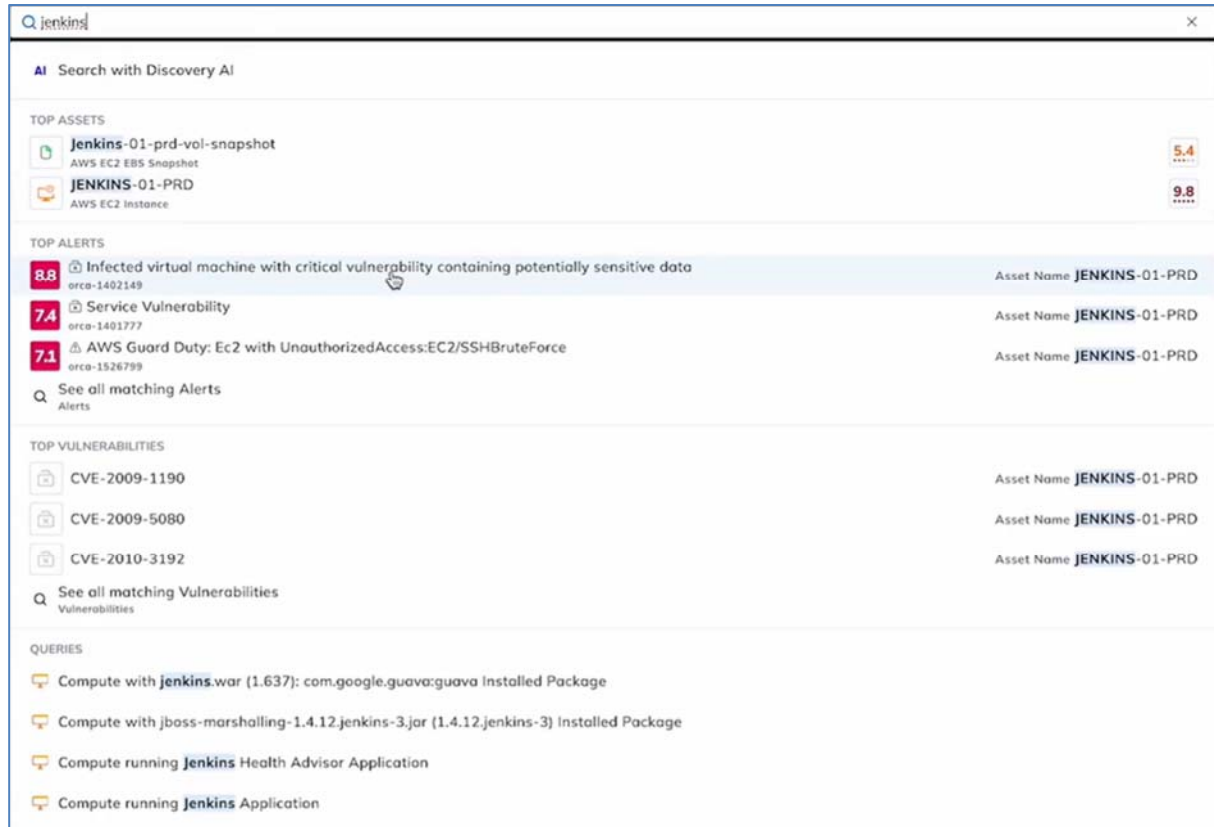
Elevate Cloud Security with Orca AI

"During incident response, seconds matter. Security operators at most firms struggle to build, let alone maintain runbooks that can keep up with the speed of business. Generative AI and LLMs offer proven relief for these teams, so they can remain focused on what matters while continuing to raise the bar on security. Orca's continued use of AI to power remediation shows how it can benefit security teams."



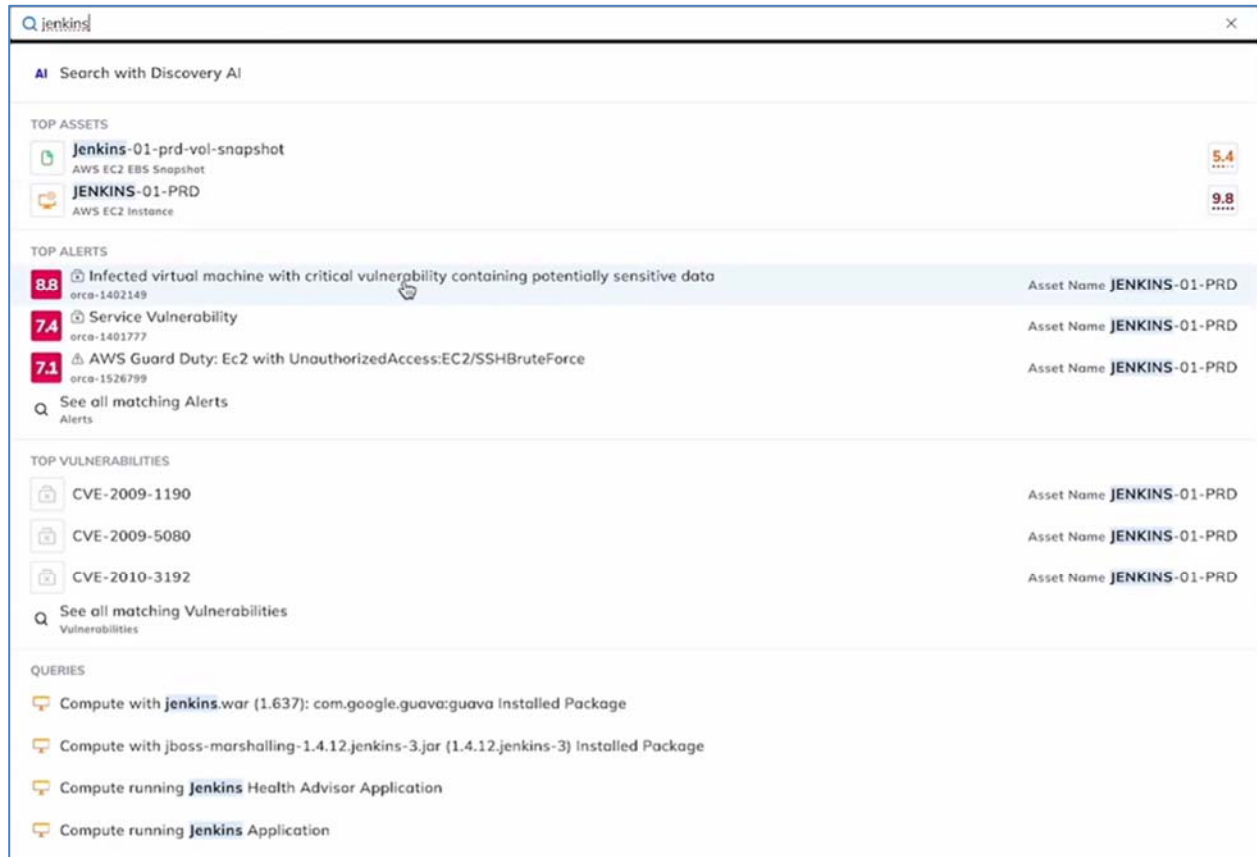
Kathy Wang
CISO and Advisor

Further, Orca provides AI powered search for a user to input a prompt based on the received alert and/or user search. *See, e.g., AI-Driven Cloud Security*, <https://orca.security/platform/ai-cloud-security/> (A user can input the search term “Jenkins” based on an alert).

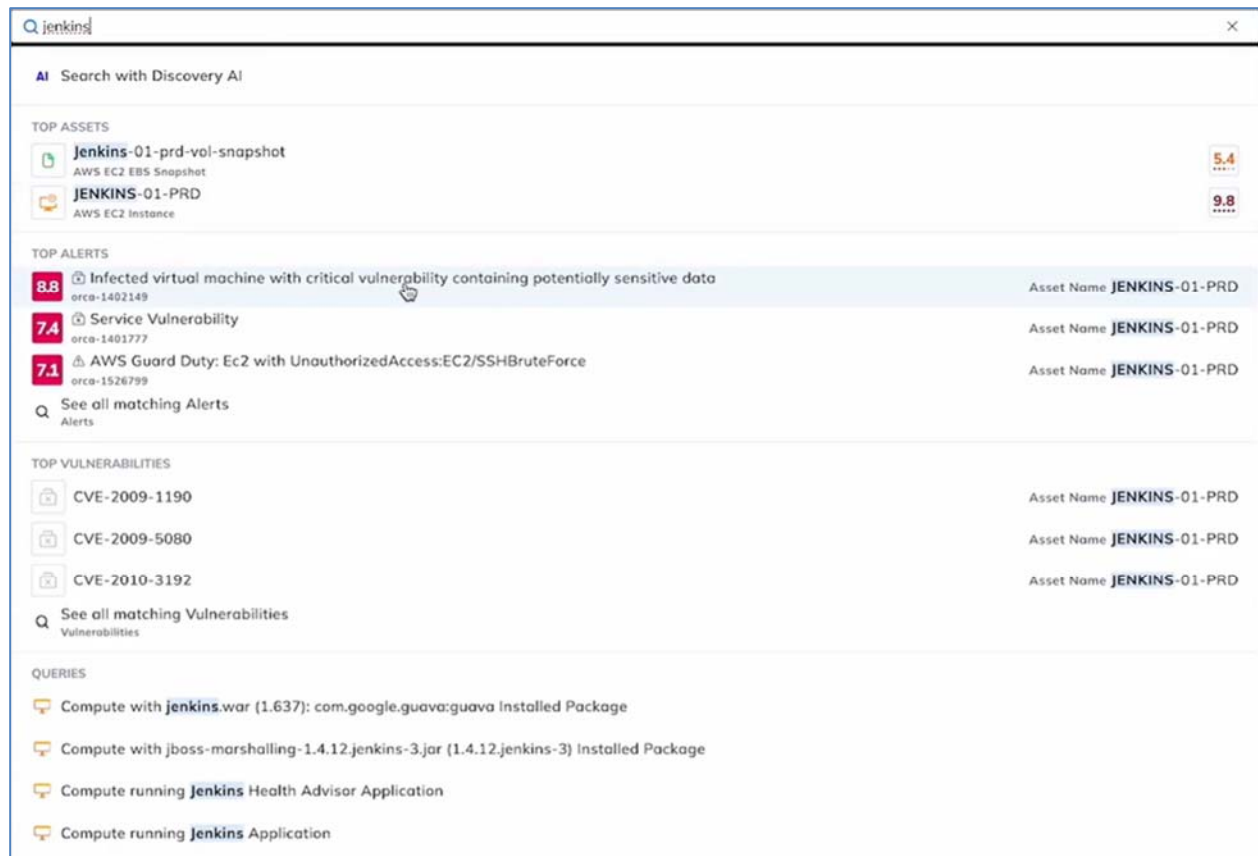


98. Claim 1 further recites “generating a prompt for a large language model (LLM) based on the received incident input” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s “Search with Discovery AI” feature. Orca provides AI powered search for a user to input a prompt based on the received alert and/or user search. *See, e.g., AI-Driven Cloud Security*, <https://orca.security/platform/ai-cloud-security/>

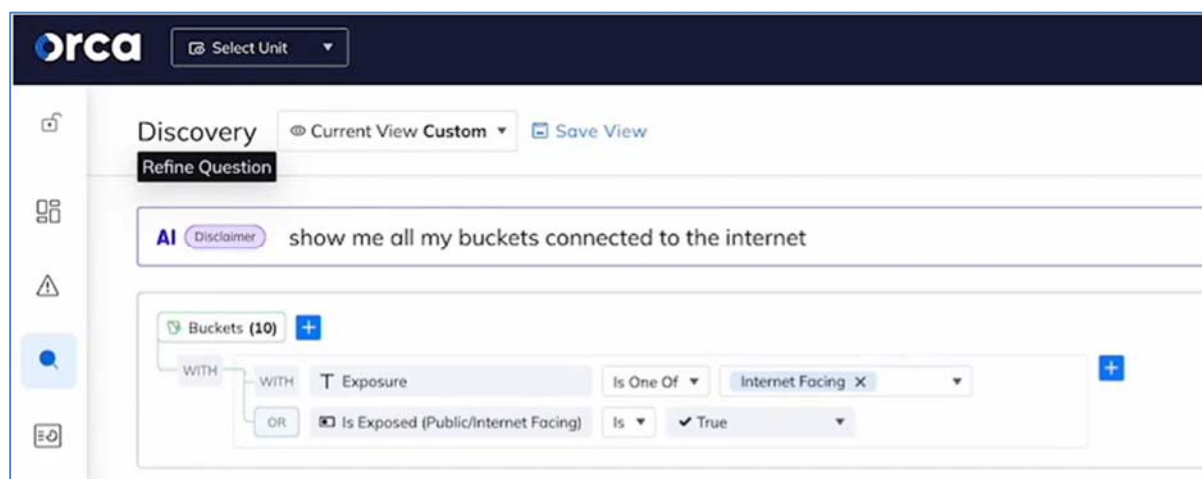
(A user can input the search term “Jenkins” based on an alert).



99. Claim 1 further recites “configuring the LLM to generate an output based on the generated prompt” Upon information and belief, Orca practices this step by, for example but not limited to, using Orca’s AI Driven Cloud Security. Using “Search with Discovery AI” to search for the term “jenkins” provides a list of top alerts, top vulnerabilities and queries. *See, e.g., AI-Driven Cloud Security, <https://orca.security/platform/ai-cloud-security/>.*

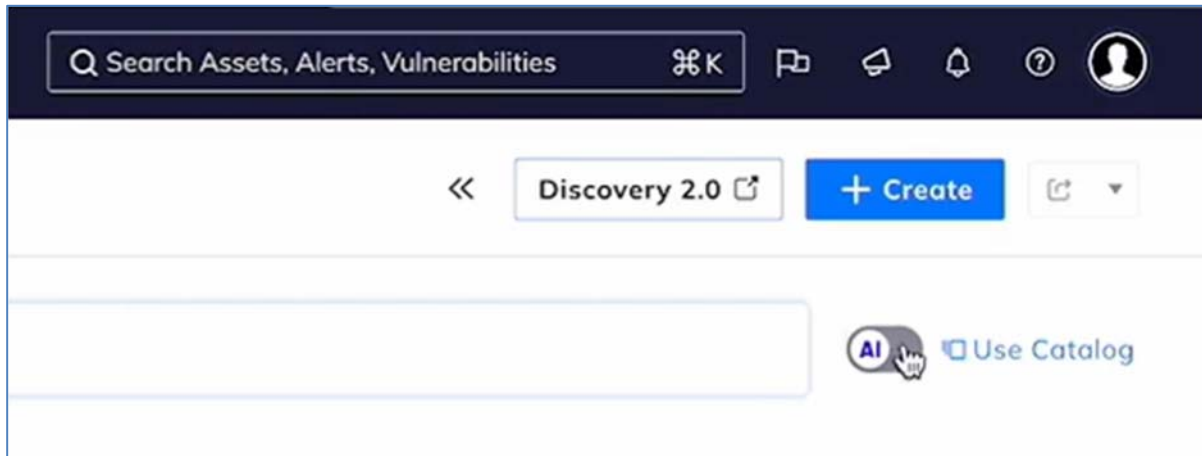


Further, Orca allows user generated prompts for its LLM through its search functionality with AI enabled. See, e.g., *AI-Driven Cloud Security*, <https://orca.security/platform/ai-cloud-security/> (“Show me all my buckets connected to the internet”).



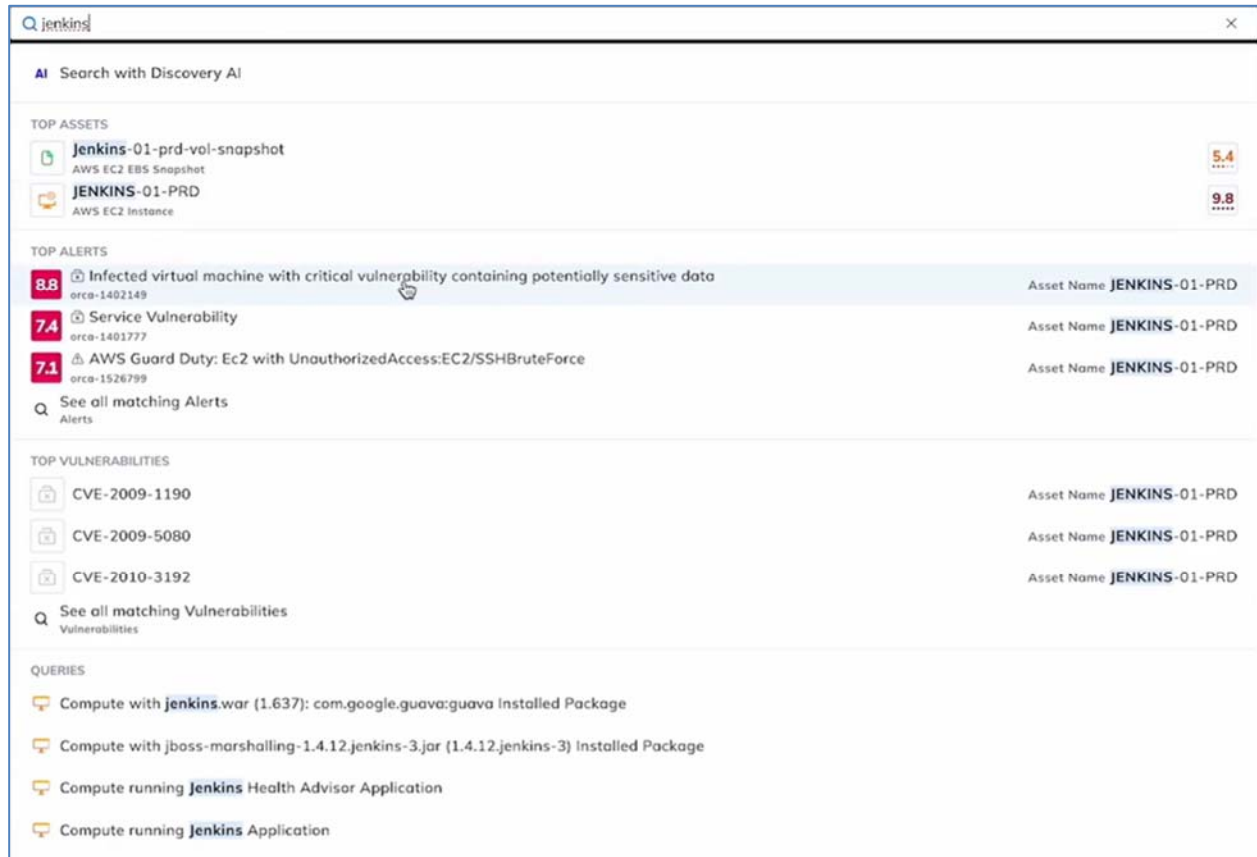
Further, Orca allows users to toggle the LLM to be configured to produce results. *See, e.g., AI-*

Driven Cloud Security, <https://orca.security/platform/ai-cloud-security/>.



100. Claim 1 further recites “mapping the received incident input into a scenario of a plurality of scenarios based on the output of the LLM, wherein each scenario is associated with an incidence response” On information and belief, Orca practices this step by, for example but not limited to, using Orca’s “Search with Discovery AI” to generate a plurality of applicable assets, alerts, vulnerabilities and queries. *See, e.g., AI-Driven Cloud Security*, <https://orca.security/platform/ai-cloud-security/>. (“I could search for Jenkins. Now, did I mean assets with the name Jenkins? Did I mean alerts that impact assets with the name Jenkins or perhaps vulnerabilities. Or maybe I meant the actual installed package. All of these things derived from that question and clicking into them I can drive further into more details to find all the different types of compute service that has an installed package which is that one there. And all done through that very simple click interface and there I have my two different servers, each

with Jenkins running on them.”).



101. Claim 1 further recites “generating a query based on the received incident input and the mapped scenario” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s “Search with Discovery AI.” Upon information and belief, when a user clicks on a listed alert, vulnerability or query Orca generates a query based on the input of “Jenkins” and the provided list of alerts. *See, e.g., AI-Driven Cloud Security*, <https://orca.security/platform/ai-cloud-security/> (“I could search for Jenkins. Now, did I mean assets with the name Jenkins? Did I mean alerts that impact assets with the name Jenkins or perhaps vulnerabilities. Or maybe I meant the actual installed package. All of these things derived from that question and clicking into them I can drive further into more details to find all the different types of compute service that has an installed package which is that one there. And

all done through that very simple click interface and there I have my two different servers, each with Jenkins running on them.”).

Q jenkins

AI Search with Discovery AI

TOP ASSETS

Jenkins-01-prd-vol-snapshot

AWS EC2 EBS Snapshot

5.4

JENKINS-01-PRD

AWS EC2 Instance

9.8

TOP ALERTS

8.8

Infected virtual machine with critical vulnerability containing potentially sensitive data

Asset Name JENKINS-01-PRD

7.4

Service Vulnerability

Asset Name JENKINS-01-PRD

7.1

AWS Guard Duty: Ec2 with UnauthorizedAccess:EC2/SSHBruteForce

Asset Name JENKINS-01-PRD

See all matching Alerts

TOP VULNERABILITIES

CVE-2009-1190

Asset Name JENKINS-01-PRD

CVE-2009-5080

Asset Name JENKINS-01-PRD

CVE-2010-3192

Asset Name JENKINS-01-PRD

See all matching Vulnerabilities

QUERIES

Compute with jenkins.war (1.637): com.google.guava:guava Installed Package

Compute with jboss-marshalling-1.4.12-jenkins-3.jar (1.4.12-jenkins-3) Installed Package

Compute running Jenkins Health Advisor Application

Compute running Jenkins Application

Query

Compute Services (89)

WITH

LINKED

Installed Package

Installed Packages

WITH

T Name

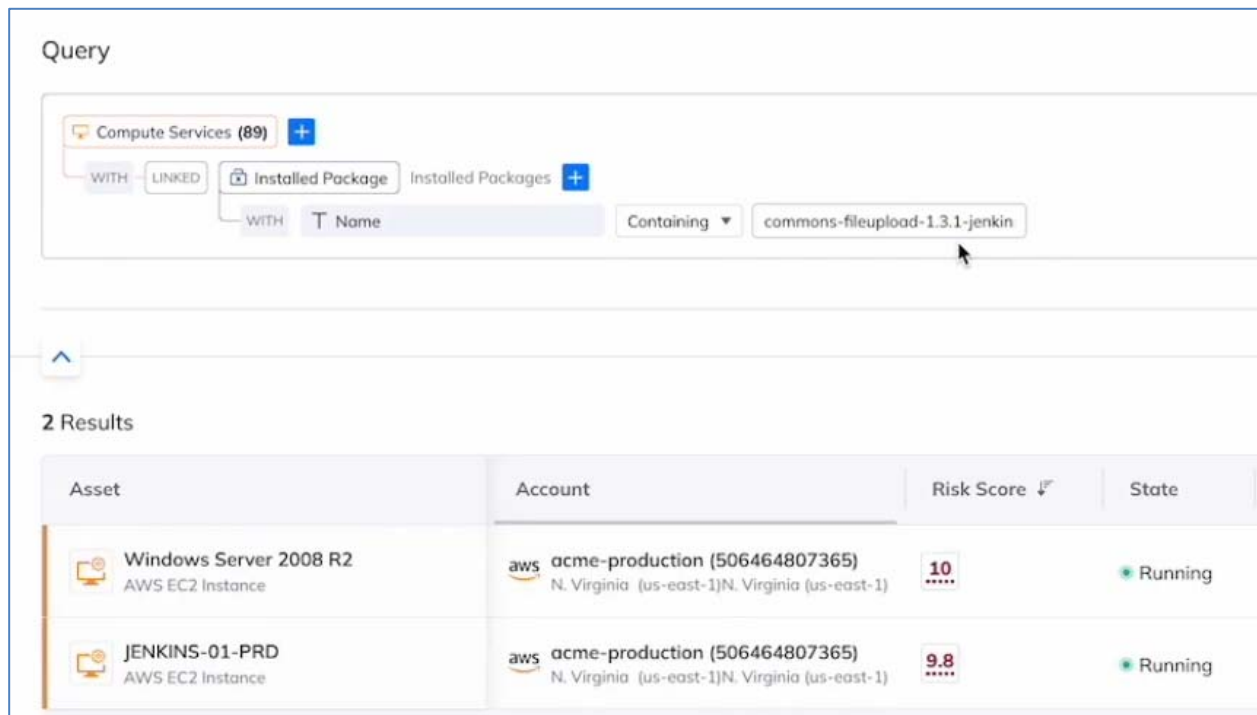
Containing

commons-fileupload-1.3.1-jenkins

2 Results

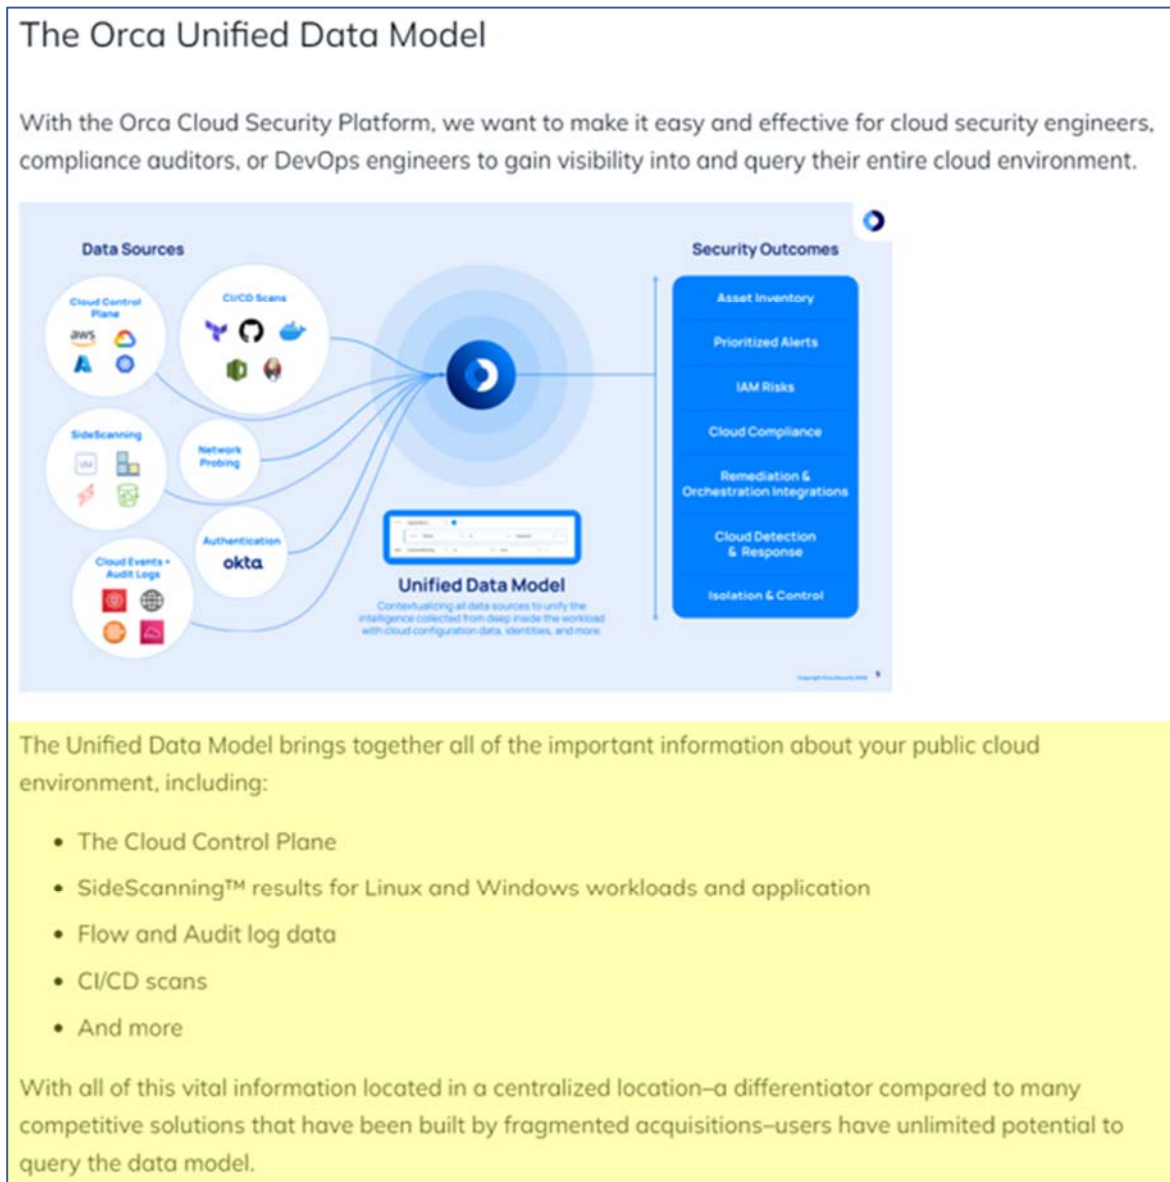
Asset	Account	Risk Score	State
<div>Windows Server 2008 R2</div> <div>AWS EC2 Instance</div>	<div>aws acme-production (506464807365)</div> <div>N. Virginia (us-east-1)N. Virginia (us-east-1)</div>	<div>10</div> <div>.....</div>	<div>Running</div>
<div>JENKINS-01-PRD</div> <div>AWS EC2 Instance</div>	<div>aws acme-production (506464807365)</div> <div>N. Virginia (us-east-1)N. Virginia (us-east-1)</div>	<div>9.8</div> <div>.....</div>	<div>Running</div>

102. Claim 1 further recites “executing the query on a security database, the security database including a representation of a computing environment; and” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s “Search with Discovery AI.” Upon information and belief, when a user clicks on a listed alert, vulnerability or query Orca executes the query on a security database including its “Unified Data Model.” *See, e.g., AI-Driven Cloud Security*, <https://orca.security/platform/ai-cloud-security/>.



103. Orca also promotes its “Unified Data Model.” *See, e.g., Cloud Security Simplified: Easily Query Your Entire Cloud Environment*, <https://orca.security/resources/blog/orca-sonar-data-query-builder/>, (“The Unified Data Model brings together all of the important information about your public cloud environment”); and its “Data Security and Posture Management.” *See e.g., Data Security and Posture Management*, <https://orca.security/platform/data-security-and-posture-management-dspm/> (“The Orca Cloud Security Platform performs continuous discovery of data stores across your cloud estate, and

alerts to security and compliance risks, without requiring any additional tools. Instead of focusing solely on data security, Orca delivers a comprehensive, context-driven picture of sensitive data exposure, enabling organizations to prioritize risks effectively, reduce alert fatigue, and stay focused on what matters most—from a single platform.”).



104. Finally, claim 1 recites “initiating a mitigation action based on a result of the executed query.” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, allowing users to take remedial steps based on the query results. *See, e.g., AI-*

Driven Cloud Security, <https://orca.security/platform/ai-cloud-security/>.

2 Results				
Asset	Account	Risk Score	State	Alerts on Asset
Windows Server 2008 R2 AWS EC2 Instance	acme-production (506464807365) N. Virginia (us-east-1)N. Virginia (us-east-1)	10 	Running	1 8 60 81 304
JENKINS-01-PRD AWS EC2 Instance	acme-production (506464807365) N. Virginia (us-east-1)N. Virginia (us-east-1)	9.8 	Running	2 6 85 62 127

Further, Orca practices this step by, for example, using Orca’s AI Driven Cloud Security to generate policies or remediation steps. *See, e.g., AI-Driven Cloud Security,* <https://orca.security/platform/ai-cloud-security/>.

Recommended Remediation

115

Delete the infected file immediately, and audit your asset and network for any anomalies or other infected files. Patch the vulnerable service according to the CVE details.

Remediation

Remediation Console:

1. Remove the infected file immediately.

2. Audit your asset and network for any anomalies or other infected files.

3. In case you are not familiar with the file and further security information is needed, you may search the file hash in [VirusTotal](#).

4. In case the file is not known to the website, you may upload it via [VirusTotal files upload](#), and the file will get scanned by various security vendors.
Important: Files that are uploaded to VirusTotal must not include proprietary information as these are publicly shared by the website.

5. Patch the vulnerable service according to the CVE details.

As a Further example, Orca’s IAM Policy Optimizer powered by Orca AI recommends remediation plans to users. *See id.*



105. Upon information and belief, Orca further infringes multiple additional claims of the '549 patent. For example, but not limited to, claims 2, 3.

106. Orca is aware of the '549 patent and Orca's infringement thereof at least as of the filing of these counterclaims. Moreover, this is another in a long line of examples of Orca copying Wiz. Accordingly, Orca has and continues to willfully infringement the '549 patent.

107. Orca has induced and continues to induce infringement of one or more claims of the '549 patent by encouraging customers to use AI-Driven Cloud Security in a manner that directly infringes those claims. Despite its knowledge of the existence of the '549 patent, since at least the filing of this Counterclaim, Orca, upon information and belief, continues to encourage, instruct, enable and otherwise cause its customers to use AI-Driven Cloud Security in a manner that infringes one or more claims of the '549 patent. Upon information and belief, Orca specifically intends that its customers use AI-Driven Cloud Security in a manner that infringes one or more claims of the '549 patent by, at a minimum, providing instructions and/or support documentation directing customers on how to use AI-Driven Cloud Security in an infringing manner, in violation of 35 U.S.C. § 271(b). For example, Orca's public web posts cited above provide instructions and encourage customers to practice all steps of the claimed method. Further, Orca provides video explanation of how to use Orca's AI-Driven Cloud Security. See <https://orca.security/platform/ai-cloud-security>.

108. Orca has contributed and continues to contribute to the infringement of one or more claims of the '549 patent. Upon information and belief, Orca knows that its AI-Driven Cloud Security features are especially made and/or adapted for users to infringe one or more claims of the '549 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use. Because Orca included features, such as for example but not limited to, AI-Driven Cloud Security in Orca's products, Orca intends for customers to use it. Upon information and belief, AI-Driven Cloud Security has no suitable use that is non-

infringing, and therefore Orca intends for customers to use AI-Driven Cloud Security in an infringing manner. Orca's sales of products including AI-Driven Cloud Security constitute contributory infringement in violation of 35 U.S.C. § 271(c).

COUNTERCLAIM V
(INFRINGEMENT OF U.S. PATENT NO. 12,003,529)

109. Wiz is the sole and exclusive owner, by assignment, of all rights, title and interest in U.S. Patent No. 12,003,549 (the "'529 patent"), entitled "Techniques for Detecting Artificial Intelligence Model Cybersecurity Risk in a Computing Environment." The '529 patent was duly and legally issued by the U.S. Patent and Trademark Office on June 4, 2024. The named inventors of the '529 patent are Amitai Cohen, Barak Sharoni, Shir Tamari, George Pisha, Itay Arbel, Daniel Velikanski, and Yaniv Shaked. A copy of the '529 patent is attached as Exhibit E.

110. The '529 patent generally relates to detecting an artificial intelligence model cybersecurity risk in a computing environment. *See* '529 patent at 1:61-64. This includes "generating a representation of the AI model in a security database[.]" *Id.* at 1:66-67. The patent further includes "initiating a mitigation action based on the cybersecurity risk." *Id.* at 2:11-12.

111. The '529 patent discloses "generating an inspectable disk based on an original disk of a resource deployed in the computing environment; and inspecting the inspectable disk for the AI model." *Id.* at 2:18-21. The patent discloses that the inspection may include "detecting an artifact of the AI model[.]" detecting "an AI model configured to execute a code object[.]" detecting "metadata of the AI model, where the metadata indicates that the AI model is a cybersecurity risk[.]" and "detecting in the AI model any one of: a secret, a certificate, a code, and any combination thereof." *Id.* at 2:21-29. "In some embodiments, a mitigation action includes revoking access from a principal, revoking access to a resource, revoking access from a resource, sandboxing a resource, revoking access to an AI model, revoking access from an AI

model, denying network communication directed to the AI model, such as network communication including a prompt for the AI model, generating an alert, updating an alert, generating an alert severity, updating an alert severity, various combinations thereof, and the like.” *Id.* at 17:24-32.

112. The ’529 patent further discloses that in response to detecting a cybersecurity risk of an AI the system “my furthermore initiate a mitigation action based on the cybersecurity risk.” *Id.* at 3:11-12. The mitigation action may be “based on a toxic combination, the detected cybersecurity object, an AI model type, a combination thereof, and the like.” *Id.* at 17:20-23.

113. Orca has infringed and continues to directly infringe one or more claims of the ’529 patent by making, using, selling, offering for sale, and/or importing into the United States without authority or license, the Orca Platform with AI Security Posture Management (AI-SPM) in violation of 35 U.S.C. § 271(a). Orca’s infringement includes infringement of, for example, claim 1 of the ’529 patent.

114. Claim 1 of the ’529 patent recites:

1. A method for detecting a cybersecurity risk of an artificial intelligence (AI), comprising:
 generating an inspectable disk based on an original disk of a resource deployed in a computing environment;
 inspecting the inspectable disk for an AI model;
 generating a representation of the AI model in a security database, the security database including a representation of the computing environment;
 inspecting the AI model for a cybersecurity risk;
 generating a representation of the cybersecurity risk in the security database, the representation of the cybersecurity risk connected to the representation of the AI model in response to detecting the cybersecurity risk; and
 initiating a mitigation action based on the cybersecurity risk.

115. On information and belief, Orca practices each and every limitation of claim 1 of the ’529 patent by and through the use of Orca’s AI Security Posture Management (“AI-SPM”).

116. The preamble of claim 1 recites “[a] method for detecting a cybersecurity risk of an artificial intelligence (AI), comprising:” To the extent the preamble is limiting, Orca practices this step by, for example but not limited to, using its AI Security Posture Management features to detect a cybersecurity risk of an artificial intelligence. *See, e.g., Orca Adds AI Security to Cloud Security Platform*, <https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/> (“Today we are pleased to announce that the Orca Cloud Security Platform now offers complete end-to-end AI Security Posture Management (AI-SPM) capabilities, so Orca customers can continue to leverage AI at unhindered speed, but do so safely.”).

Organizations are increasingly leveraging Generative AI and Large Language Models (LLMs) to optimize business processes and improve products and services. A [recent Gartner report](#) predicts that global spending on AI software will grow to \$298 billion by 2027, with a compound annual growth rate (CAGR) of 19.1%. Some may even be of the opinion that these figures are conservative.

From assets scanned by the Orca Cloud Security Platform, we found that over 37% of organizations have already adopted at least one AI service, with the highest use being Amazon SageMaker and Bedrock (68%) followed by Azure OpenAI (50%), and Vertex AI comes in third at 21%.

Along with all the great business benefits and improved products and services that AI brings, there is however one catch: security. In our [2024 State of Cloud Security Report](#), we found that 82% of Amazon SageMaker users have exposed notebooks, meaning that they are publicly accessible. Especially since AI models often include sensitive data and intellectual property in their training data, these cloud resources are at even higher potential risk than other resources. This has given rise to a new type of security need: AI Security.

While the exposure of SageMaker notebooks is a significant concern, it is a condition that falls within the user’s responsibility under the shared responsibility model. Therefore, it’s important that organizations follow best practices and use the right AI security tools so they can easily identify and remediate these exposures, ensuring the secure deployment of their AI and ML workloads.

Today we are pleased to announce that the [Orca Cloud Security Platform](#) now offers complete end-to-end AI Security Posture Management (AI-SPM) capabilities, so Orca customers can continue to leverage AI at unhindered speed, but do so safely. By offering AI security from Orca’s comprehensive platform, organizations avoid having to add yet another point security solution specific to AI security to their arsenal, reducing overhead and integrating into existing workflows.

117. Claim 1 further recites “generating an inspectable disk based on an original disk of a resource deployed in a computing environment” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s SideScanning Technology “to cover AI models” in combination with Orca’s AI-SPM. *See, e.g., Orca Adds AI Security to Cloud Security Platform*, <https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/> (“Leveraging Orca’s patented agentless SideScanning technology, we’ve extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources . . . Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM).”).

Why is Orca adding AI-SPM?

Many of the security risks facing AI models and LLMs are similar to other cloud assets, including limited visibility, accidental public access, unencrypted sensitive data, shadow data, and unsecured keys.

Leveraging Orca's patented agentless [SideScanning technology](#), we've extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources. In addition, we're applying our existing technology for use cases that are unique to AI security, such as detecting sensitive data in training sets, that could later be unintentionally exposed by the LLM or Generative AI application.

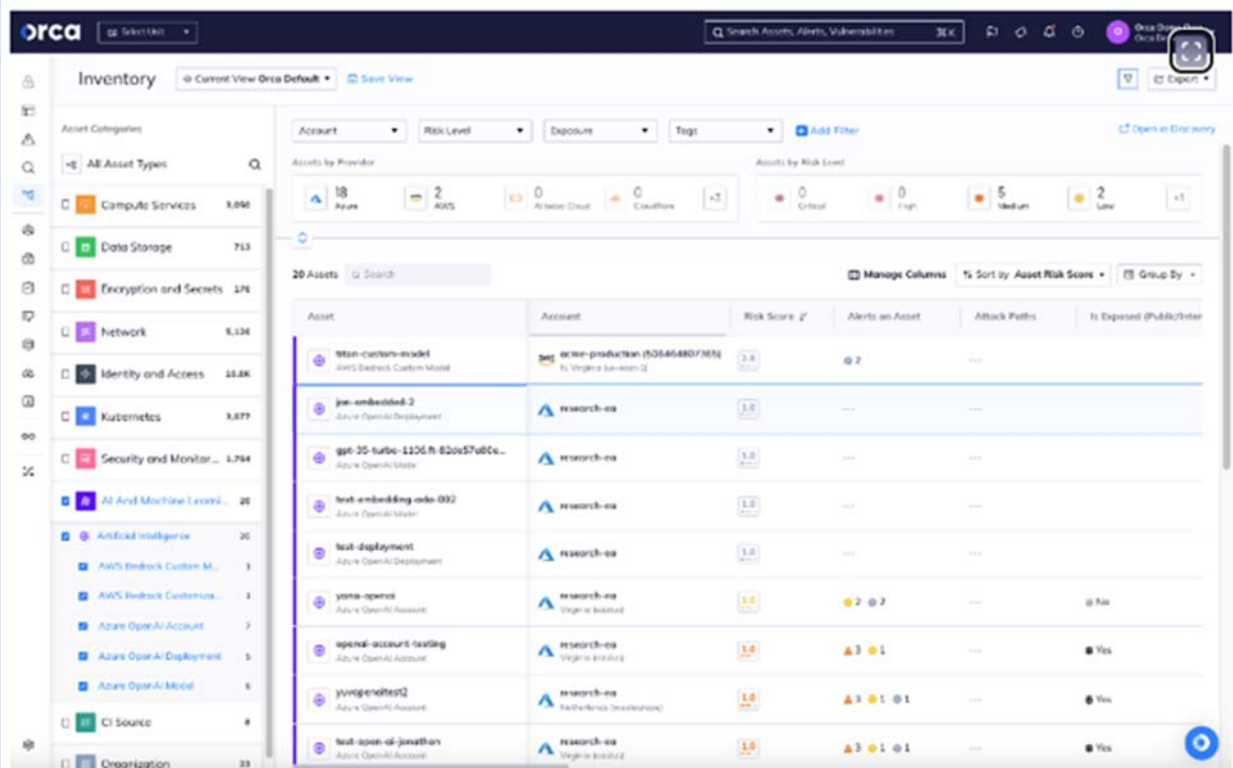
Since the Orca platform does not require agents to inspect cloud resources, coverage is always continuous and 100%, and will immediately scan any new AI resources as soon as they are deployed and alert to any detected risks. However Orca goes beyond just covering AI models – the platform also maps out all the AI related packages that are used to train or infer AI.

By adding AI security to our comprehensive cloud security platform, organizations don't need to procure, deploy, or learn how to use another separate point solution but can instead leverage one unified platform for all cloud security needs.

#1. AI and ML BOM + inventory

Challenge: Much like other resources in the cloud, shadow AI and LLMs are a major concern. In their excitement to explore all the opportunities that Generative AI brings, developers are not always waiting for IT approval before integrating Gen AI services into their flows. And security is probably not the first thing on their mind. This leaves security teams in the dark about which AI projects have been deployed in their environment, whether they contain sensitive data, and whether they are secure.

Orca solution: Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM). Orca detects projects on Azure OpenAI, Amazon Bedrock, Google Vertex AI, AWS Sagemaker and those using one of 50+ most commonly used AI software packages, including Pytorch, TensorFlow, OpenAI, Hugging Face, scikit-learn, and many more.



The Orca Platform shows a full inventory and bill of materials of all deployed AI models

118. Claim 1 further recites “inspecting the inspectable disk for an AI model”

Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s SideScanning Technology “to cover AI models” in combination with Orca’s AI-SPM. *See, e.g., Orca Adds AI Security to Cloud Security Platform,*

<https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/>

(“Leveraging Orca’s patented agentless SideScanning technology, we’ve extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources . . . Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM).”).

Why is Orca adding AI-SPM?

Many of the security risks facing AI models and LLMs are similar to other cloud assets, including limited visibility, accidental public access, unencrypted sensitive data, shadow data, and unsecured keys. Leveraging Orca's patented agentless [SideScanning technology](#), we've extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources. In addition, we're applying our existing technology for use cases that are unique to AI security, such as detecting sensitive data in training sets, that could later be unintentionally exposed by the LLM or Generative AI application.

Since the Orca platform does not require agents to inspect cloud resources, coverage is always continuous and 100%, and will immediately scan any new AI resources as soon as they are deployed and alert to any detected risks. However Orca goes beyond just covering AI models – the platform also maps out all the AI related packages that are used to train or infer AI.

By adding AI security to our comprehensive cloud security platform, organizations don't need to procure, deploy, or learn how to use another separate point solution but can instead leverage one unified platform for all cloud security needs.

#1. AI and ML BOM + inventory

Challenge: Much like other resources in the cloud, shadow AI and LLMs are a major concern. In their excitement to explore all the opportunities that Generative AI brings, developers are not always waiting for IT approval before integrating Gen AI services into their flows. And security is probably not the first thing on their mind. This leaves security teams in the dark about which AI projects have been deployed in their environment, whether they contain sensitive data, and whether they are secure.

Orca solution: Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM). Orca detects projects on Azure OpenAI, Amazon Bedrock, Google Vertex AI, AWS Sagemaker and those using one of 50+ most commonly used AI software packages, including Pytorch, TensorFlow, OpenAI, Hugging Face, scikit-learn, and many more.

Asset	Account	Risk Score	Alerts on Asset	Attack Paths	Is Exposed (Public/Inter)
llm-custom-model AWS Bedrock Custom Model	acme-production (303464807389) N. Virginia (us-east-2)	2.8	2	---	---
jai-embedding-2 Azure OpenAI Deployment	research-na	1.0	---	---	---
gpt-35-turbo-1106-hf-02057ulb... Azure OpenAI Model	research-na	1.0	---	---	---
text-embedding-ada-002 Azure OpenAI Model	research-na	1.0	---	---	---
text-deployment Azure OpenAI Deployment	research-na	1.0	---	---	---
yama-openai Azure OpenAI Account	research-na Virginia (us-east-2)	1.0	2 2	---	at Risk
openai-account-testing Azure OpenAI Account	research-na Virginia (us-east-2)	1.0	3 1	---	Yes
yungendites2 Azure OpenAI Account	research-na Texas (us-central-2)	1.0	1 1 1	---	Yes
text-openai-jurathion Azure OpenAI Account	research-na Virginia (us-east-2)	1.0	3 1 1	---	Yes

The Orca Platform shows a full inventory and bill of materials of all deployed AI models

119. Claim 1 further recites “generating a representation of the AI model in a security database, the security database including a representation of the computing environment” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s SideScanning Technology “to cover AI models” in combination with Orca’s AI-

SPM. See, e.g., *Orca Adds AI Security to Cloud Security Platform*,

<https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/>

(“Leveraging Orca’s patented agentless SideScanning technology, we’ve extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources... Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM).”).

Why is Orca adding AI-SPM?

Many of the security risks facing AI models and LLMs are similar to other cloud assets, including limited visibility, accidental public access, unencrypted sensitive data, shadow data, and unsecured keys. Leveraging Orca’s patented agentless [SideScanning technology](#), we’ve extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources. In addition, we’re applying our existing technology for use cases that are unique to AI security, such as detecting sensitive data in training sets, that could later be unintentionally exposed by the LLM or Generative AI application.

Since the Orca platform does not require agents to inspect cloud resources, coverage is always continuous and 100%, and will immediately scan any new AI resources as soon as they are deployed and alert to any detected risks. However Orca goes beyond just covering AI models – the platform also maps out all the AI related packages that are used to train or infer AI.

By adding AI security to our comprehensive cloud security platform, organizations don’t need to procure, deploy, or learn how to use another separate point solution but can instead leverage one unified platform for all cloud security needs.

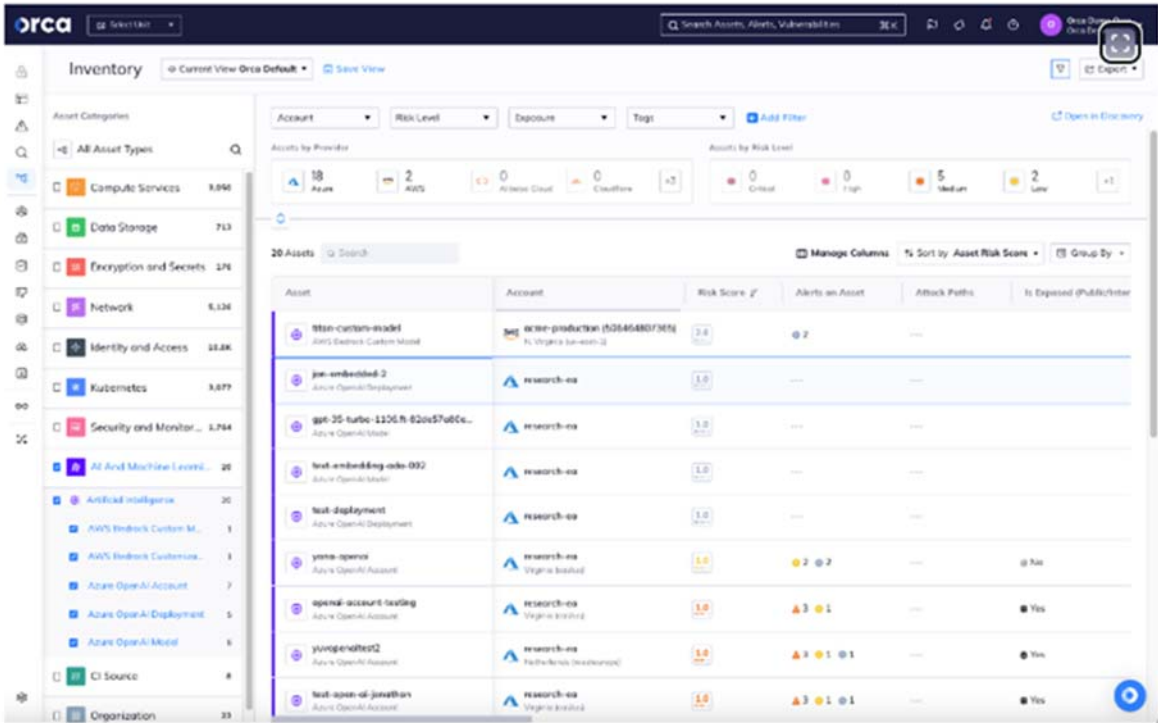
What are Orca's AI Security capabilities?

Orca includes a new AI Security dashboard that provides an overview of the AI models that are deployed in the cloud environment, what data they contain, and whether they are at risk. Orca covers risks end-to-end, from training and fine-tuning to production deployment and inference. Orca connects with your AI data and processes on all levels – the managed cloud AI services, the unmanaged AI models and packages that your developers are using, and even with shift-left detection of leaky secret tokens AI services in your codebase.

#1. AI and ML BOM + inventory

Challenge: Much like other resources in the cloud, shadow AI and LLMs are a major concern. In their excitement to explore all the opportunities that Generative AI brings, developers are not always waiting for IT approval before integrating Gen AI services into their flows. And security is probably not the first thing on their mind. This leaves security teams in the dark about which AI projects have been deployed in their environment, whether they contain sensitive data, and whether they are secure.

Orca solution: Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM). Orca detects projects on Azure OpenAI, Amazon Bedrock, Google Vertex AI, AWS Sagemaker and those using one of 50+ most commonly used AI software packages, including Pytorch, TensorFlow, OpenAI, Hugging Face, scikit-learn, and many more.




The Orca Platform shows a full inventory and bill of materials of all deployed AI models

120. Further, Orca practices this step by generating a representation of the AI model in its “Unified Data Model.” *See, e.g., Cloud Security Simplified: Easily Query Your Entire Cloud Environment*, <https://orca.security/resources/blog/orca-sonar-data-query-builder/> (“The Unified Data Model brings together all of the important information about your public cloud environment . . .”); and through the use of its “Data Security and Posture Management.” *See e.g., Data Security and Posture Management*, <https://orca.security/platform/data-security-and-posture-management-dspm/> (“The Orca Cloud Security Platform performs continuous discovery of data stores across your cloud estate, and alerts to security and compliance risks, without requiring any additional tools. Instead of focusing solely on data security, Orca delivers a comprehensive, context-driven picture of sensitive data exposure, enabling organizations to prioritize risks effectively, reduce alert fatigue, and stay focused on what matters most—from a single platform.”).

The Orca Unified Data Model

With the Orca Cloud Security Platform, we want to make it easy and effective for cloud security engineers, compliance auditors, or DevOps engineers to gain visibility into and query their entire cloud environment.



The Unified Data Model brings together all of the important information about your public cloud environment, including:

- The Cloud Control Plane
- SideScanning™ results for Linux and Windows workloads and application
- Flow and Audit log data
- CI/CD scans
- And more

With all of this vital information located in a centralized location—a differentiator compared to many competitive solutions that have been built by fragmented acquisitions—users have unlimited potential to query the data model.

OUR APPROACH

Detect and Prioritize Cloud Data Security Risks with Context

The Orca Cloud Security Platform performs continuous discovery of data stores across your cloud estate, and alerts to security and compliance risks, without requiring any additional tools. Instead of focusing solely on data security, Orca delivers a comprehensive, context-driven picture of sensitive data exposure, enabling organizations to prioritize risks effectively, reduce alert fatigue, and stay focused on what matters most—from a single platform.

121. Claim 1 further recites “inspecting the AI model for a cybersecurity risk”

Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s SideScanning Technology “to cover AI models” in combination with Orca’s AI-SPM. *See, e.g., Orca Adds AI Security to Cloud Security Platform*,

<https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/>

(“Leveraging Orca’s patented agentless SideScanning technology, we’ve extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources . . . Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM) . . . Since it could be very damaging if this information got reached, it is very important to ensure that AI models are not publicly exposed. However, misconfigurations do happen and especially with Shadow AI, security may not be the first thing developers have on their mind.”).

Why is Orca adding AI-SPM?

Many of the security risks facing AI models and LLMs are similar to other cloud assets, including limited visibility, accidental public access, unencrypted sensitive data, shadow data, and unsecured keys. Leveraging Orca's patented agentless [SideScanning technology](#), we've extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources. In addition, we're applying our existing technology for use cases that are unique to AI security, such as detecting sensitive data in training sets, that could later be unintentionally exposed by the LLM or Generative AI application.

Since the Orca platform does not require agents to inspect cloud resources, coverage is always continuous and 100%, and will immediately scan any new AI resources as soon as they are deployed and alert to any detected risks. However Orca goes beyond just covering AI models – the platform also maps out all the AI related packages that are used to train or infer AI.

By adding AI security to our comprehensive cloud security platform, organizations don't need to procure, deploy, or learn how to use another separate point solution but can instead leverage one unified platform for all cloud security needs.

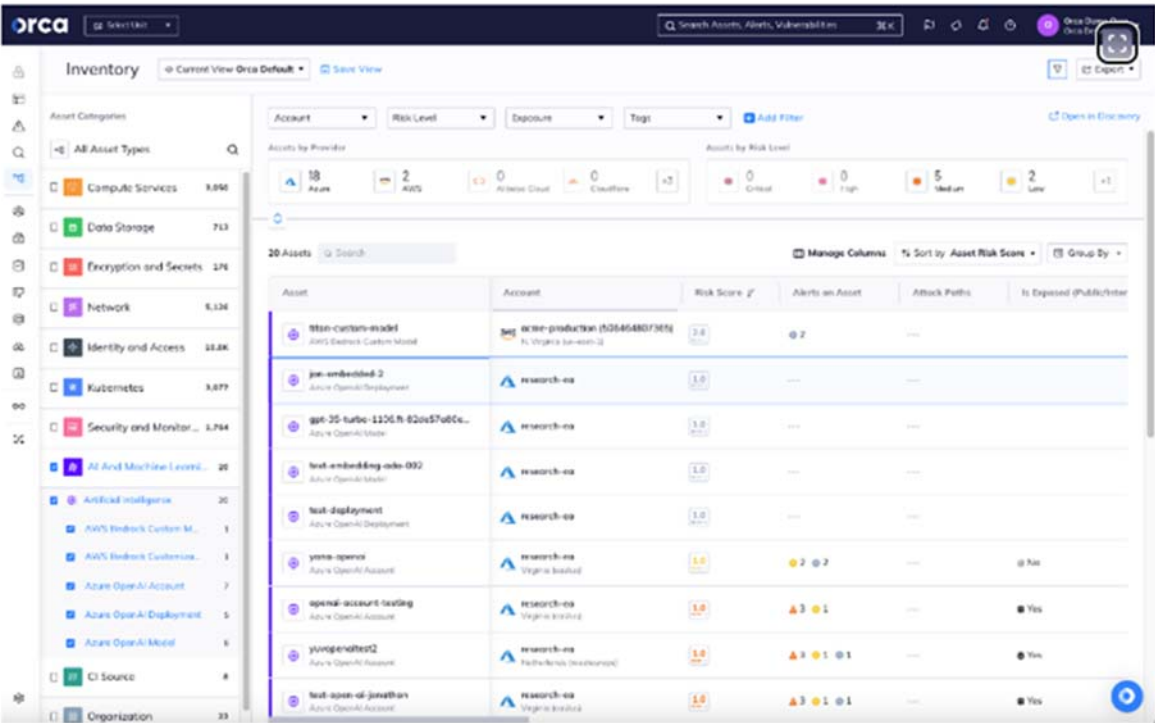
What are Orca's AI Security capabilities?

Orca includes a new AI Security dashboard that provides an overview of the AI models that are deployed in the cloud environment, what data they contain, and whether they are at risk. Orca covers risks end-to-end, from training and fine-tuning to production deployment and inference. Orca connects with your AI data and processes on all levels – the managed cloud AI services, the unmanaged AI models and packages that your developers are using, and even with shift-left detection of leaky secret tokens AI services in your codebase.

#1. AI and ML BOM + inventory

Challenge: Much like other resources in the cloud, shadow AI and LLMs are a major concern. In their excitement to explore all the opportunities that Generative AI brings, developers are not always waiting for IT approval before integrating Gen AI services into their flows. And security is probably not the first thing on their mind. This leaves security teams in the dark about which AI projects have been deployed in their environment, whether they contain sensitive data, and whether they are secure.

Orca solution: Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM). Orca detects projects on Azure OpenAI, Amazon Bedrock, Google Vertex AI, AWS Sagemaker and those using one of 50+ most commonly used AI software packages, including Pytorch, TensorFlow, OpenAI, Hugging Face, scikit-learn, and many more.

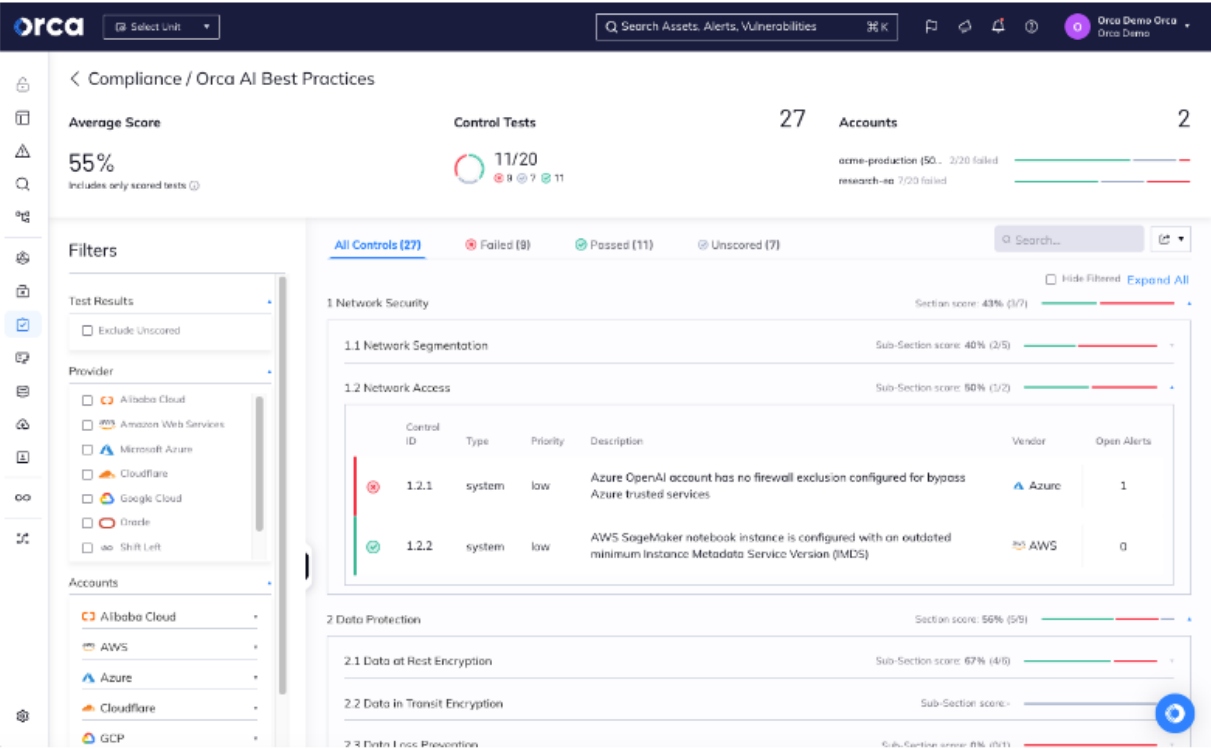


The Orca Platform shows a full inventory and bill of materials of all deployed AI models

#2. Warn when AI projects are publicly accessible

Challenge: Data used to train AI models is often sensitive and can contain proprietary information. Since it could be very damaging if this information got breached, it's very important to ensure that AI models are not publicly exposed. However, misconfigurations do happen, and especially with Shadow AI, security may not be the first thing developers have on their mind. This leaves security teams struggling to ensure that all AI models are being kept private.

Orca solution: Since Orca has insight into the AI model settings and network access, Orca will alert whenever public access is allowed, so security teams can quickly fix the issue to prevent any data breaches.

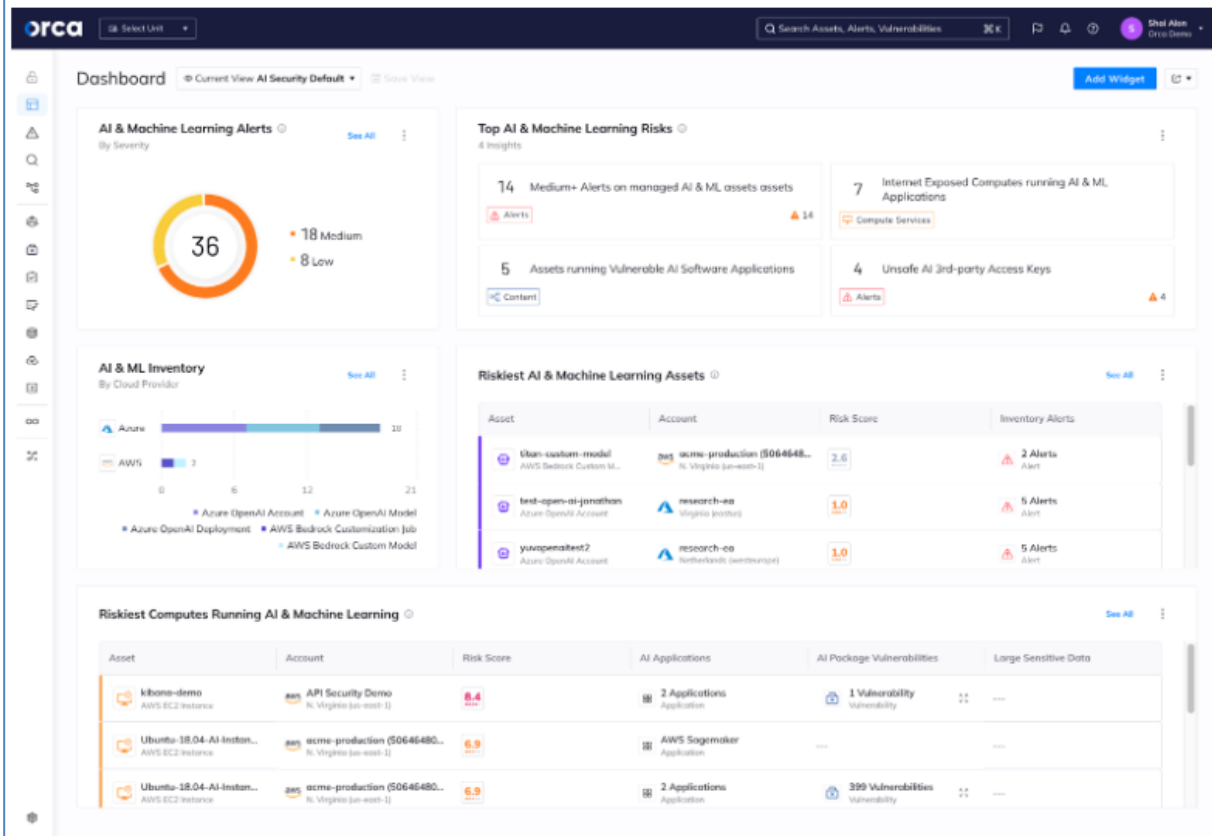


Orca AI Security best practices compliance performs AI-Security Posture Management

#3. Detect sensitive data in AI projects

Challenge: AI models use vast amounts of data to train models. It's possible that this data inadvertently contains sensitive data, and that developers and security teams are not even aware of the presence of the data, such as PII and access keys. LLMs and Generative AI models could then “spill out” this sensitive data, which could be used by bad actors. In addition, if developers are not aware that the training set contains sensitive data, they may not prioritize security of the resource as much as it needs to be.

Orca solution: Using Orca's [Data Security Posture Management \(DSPM\)](#) capabilities, Orca scans and classifies all the data stored in AI projects, and alerts if they contain sensitive data, such as email addresses, telephone numbers, email addresses, and social security numbers (PII), or personal health information (PHI). By informing security teams where sensitive data is located, they can make sure that these assets are protected with the highest level of security.



Orca displays pertinent security information in the AI Security dashboard

122. Claim 1 further recites “generating a representation of the cybersecurity risk in the security database, the representation of the cybersecurity risk connected to the representation of the AI model in response to detecting the cybersecurity risk; and . . .” Orca’s public web posts

confirm that Orca practices this step by, for example but not limited to, using Orca's SideScanning Technology "to cover AI models" in combination with Orca's AI-SPM. Orca further displays a representation of the cybersecurity risk on the AI Security Dashboard including alerting on publicly accessible AI models as well as unsecure AI training data containing sensitive data such as personal addresses or social security information. *See, e.g., Orca Adds AI Security to Cloud Security Platform*, <https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/> ("Leveraging Orca's patented agentless SideScanning technology, we've extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources . . . Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM) . . . Since Orca has insight into the AI model settings and network access, Orca will alert whenever public access is allowed, so security teams can quickly fix the issue to prevent any data breaches . . . Orca scans and classifies all the data stored in AI projects, and alerts if they contain sensitive data, such as email addresses, telephone numbers, email addresses, and social security numbers (PII), or personal health information (PHI).").

Why is Orca adding AI-SPM?

Many of the security risks facing AI models and LLMs are similar to other cloud assets, including limited visibility, accidental public access, unencrypted sensitive data, shadow data, and unsecured keys. Leveraging Orca's patented agentless [SideScanning technology](#), we've extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources. In addition, we're applying our existing technology for use cases that are unique to AI security, such as detecting sensitive data in training sets, that could later be unintentionally exposed by the LLM or Generative AI application.

Since the Orca platform does not require agents to inspect cloud resources, coverage is always continuous and 100%, and will immediately scan any new AI resources as soon as they are deployed and alert to any detected risks. However Orca goes beyond just covering AI models – the platform also maps out all the AI related packages that are used to train or infer AI.

By adding AI security to our comprehensive cloud security platform, organizations don't need to procure, deploy, or learn how to use another separate point solution but can instead leverage one unified platform for all cloud security needs.

What are Orca's AI Security capabilities?

Orca includes a new AI Security dashboard that provides an overview of the AI models that are deployed in the cloud environment, what data they contain, and whether they are at risk. Orca covers risks end-to-end, from training and fine-tuning to production deployment and inference. Orca connects with your AI data and processes on all levels – the managed cloud AI services, the unmanaged AI models and packages that your developers are using, and even with shift-left detection of leaky secret tokens AI services in your codebase.

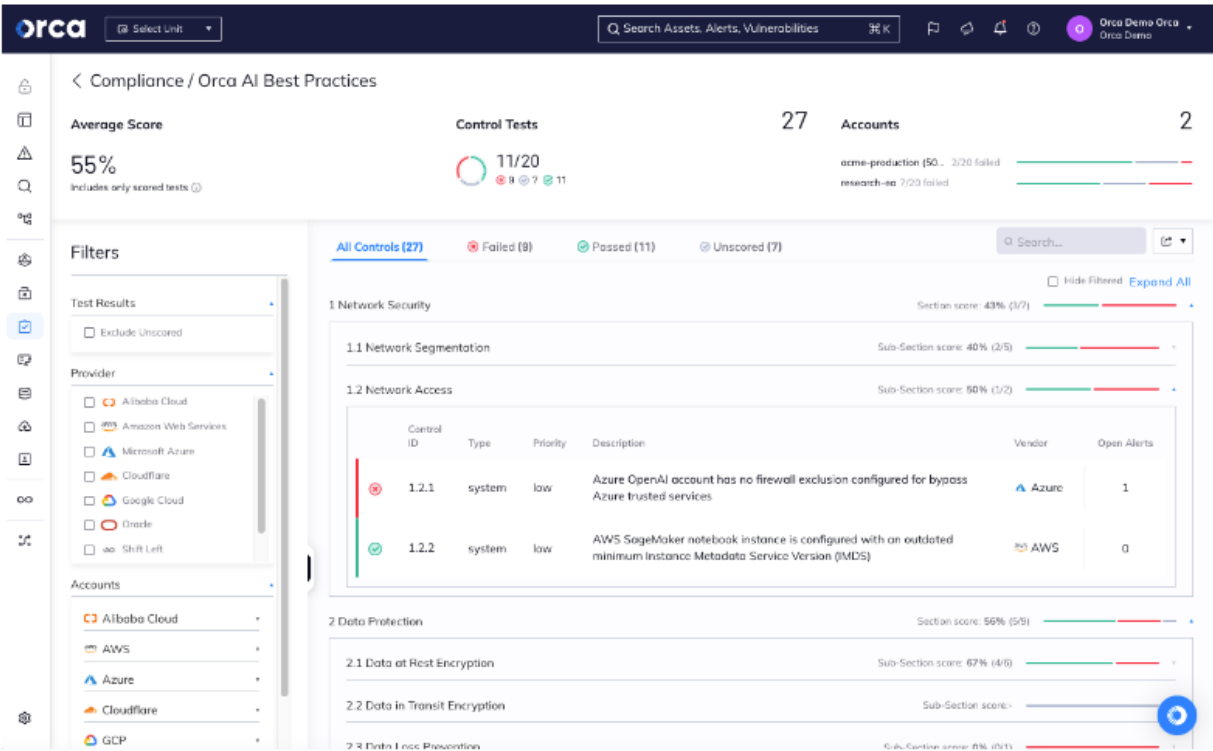
Orca solution: Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM). Orca detects projects on Azure OpenAI, Amazon Bedrock, Google Vertex AI, AWS Sagemaker and those using one of 50+ most commonly used AI software packages, including Pytorch, TensorFlow, OpenAI, Hugging Face, scikit-learn, and many more.



#2. Warn when AI projects are publicly accessible

Challenge: Data used to train AI models is often sensitive and can contain proprietary information. Since it could be very damaging if this information got breached, it's very important to ensure that AI models are not publicly exposed. However, misconfigurations do happen, and especially with Shadow AI, security may not be the first thing developers have on their mind. This leaves security teams struggling to ensure that all AI models are being kept private.

Orca solution: Since Orca has insight into the AI model settings and network access, Orca will alert whenever public access is allowed, so security teams can quickly fix the issue to prevent any data breaches.

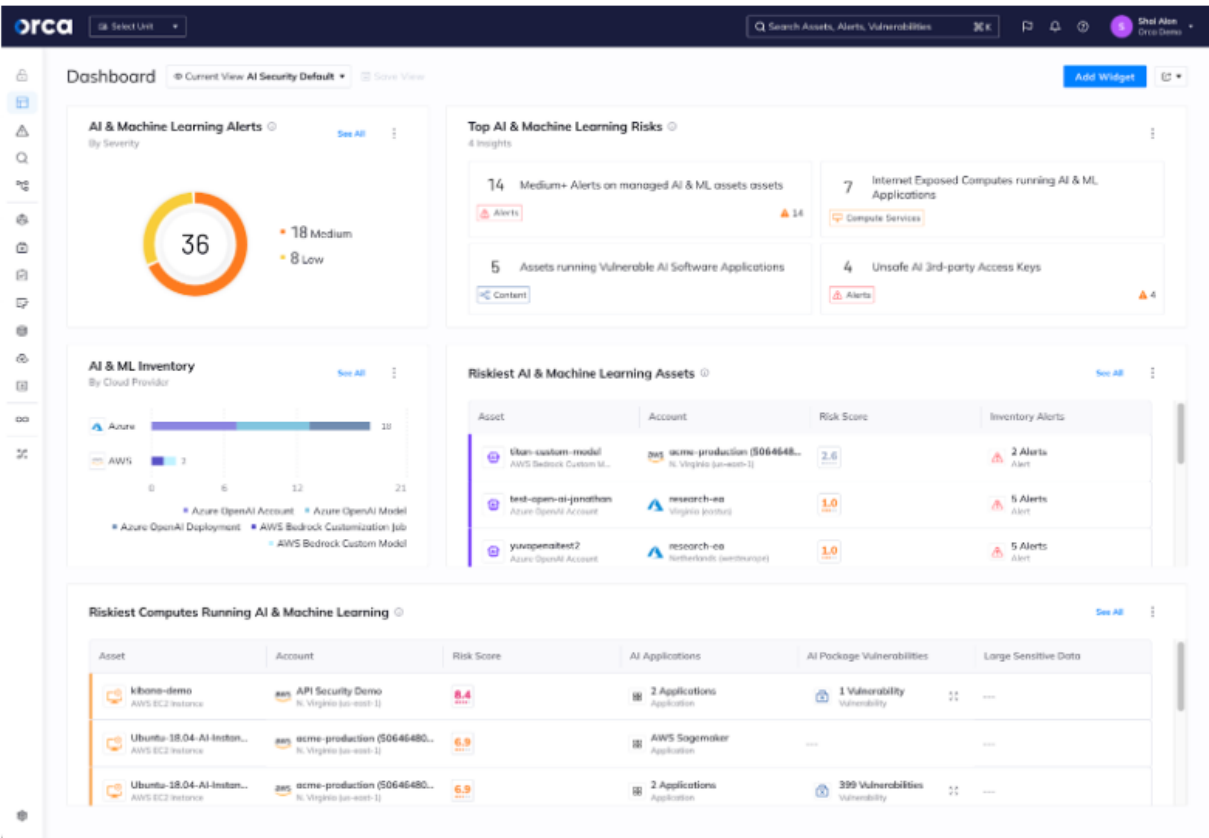


Orca AI Security best practices compliance performs AI-Security Posture Management

#3. Detect sensitive data in AI projects

Challenge: AI models use vast amounts of data to train models. It's possible that this data inadvertently contains sensitive data, and that developers and security teams are not even aware of the presence of the data, such as PII and access keys. LLMs and Generative AI models could then “spill out” this sensitive data, which could be used by bad actors. In addition, if developers are not aware that the training set contains sensitive data, they may not prioritize security of the resource as much as it needs to be.

Orca solution: Using Orca’s [Data Security Posture Management \(DSPM\)](#) capabilities, Orca scans and classifies all the data stored in AI projects, and alerts if they contain sensitive data, such as email addresses, telephone numbers, email addresses, and social security numbers (PII), or personal health information (PHI). By informing security teams where sensitive data is located, they can make sure that these assets are protected with the highest level of security.




Orca displays pertinent security information in the AI Security dashboard

123. Further, Orca practices this limitation through the use of its “Unified Data Model,” see, e.g., *Cloud Security Simplified: Easily Query Your Entire Cloud Environment*, <https://orca.security/resources/blog/orca-sonar-data-query-builder/> (“The Unified Data Model

brings together all of the important information about your public cloud environment”); and its “Data Security and Posture Management.” *See e.g., Data Security and Posture Management*, <https://orca.security/platform/data-security-and-posture-management-dspm/> (“The Orca Cloud Security Platform performs continuous discovery of data stores across your cloud estate, and alerts to security and compliance risks, without requiring any additional tools. Instead of focusing solely on data security, Orca delivers a comprehensive, context-driven picture of sensitive data exposure, enabling organizations to prioritize risks effectively, reduce alert fatigue, and stay focused on what matters most—from a single platform.”).

The Orca Unified Data Model

With the Orca Cloud Security Platform, we want to make it easy and effective for cloud security engineers, compliance auditors, or DevOps engineers to gain visibility into and query their entire cloud environment.



The Unified Data Model brings together all of the important information about your public cloud environment, including:

- The Cloud Control Plane
- SideScanning™ results for Linux and Windows workloads and application
- Flow and Audit log data
- CI/CD scans
- And more

With all of this vital information located in a centralized location—a differentiator compared to many competitive solutions that have been built by fragmented acquisitions—users have unlimited potential to query the data model.

OUR APPROACH

Detect and Prioritize Cloud Data Security Risks with Context

The Orca Cloud Security Platform performs continuous discovery of data stores across your cloud estate, and alerts to security and compliance risks, without requiring any additional tools. Instead of focusing solely on data security, Orca delivers a comprehensive, context-driven picture of sensitive data exposure, enabling organizations to prioritize risks effectively, reduce alert fatigue, and stay focused on what matters most—from a single platform.

124. Finally, claim 1 recites “initiating a mitigation action based on the cybersecurity risk.” Orca’s public web posts confirm that Orca practices this step by, for example but not limited to, using Orca’s AI-SPM to generate alerts. Orca further displays alerts on the AI Security Dashboard. *See, e.g., Orca Adds AI Security to Cloud Security Platform*, <https://orca.security/resources/blog/orca-adds-ai-security-to-cloud-security-platform/> (“Leveraging Orca’s patented agentless SideScanning technology, we’ve extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources . . . Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM) . . . Since it could be very damaging if this information got reached, it is very important to ensure that AI models are not publicly exposed. However, misconfigurations do happen and especially with Shadow AI, security may not be the first thing developers have on their mind. ”).

Why is Orca adding AI-SPM?

Many of the security risks facing AI models and LLMs are similar to other cloud assets, including limited visibility, accidental public access, unencrypted sensitive data, shadow data, and unsecured keys. Leveraging Orca's patented agentless [SideScanning technology](#), we've extended our platform to also cover AI models, providing the same risk insights and deep data that we provide on other cloud resources. In addition, we're applying our existing technology for use cases that are unique to AI security, such as detecting sensitive data in training sets, that could later be unintentionally exposed by the LLM or Generative AI application.

Since the Orca platform does not require agents to inspect cloud resources, coverage is always continuous and 100%, and will immediately scan any new AI resources as soon as they are deployed and alert to any detected risks. However Orca goes beyond just covering AI models – the platform also maps out all the AI related packages that are used to train or infer AI.

By adding AI security to our comprehensive cloud security platform, organizations don't need to procure, deploy, or learn how to use another separate point solution but can instead leverage one unified platform for all cloud security needs.


What are Orca's AI Security capabilities?

Orca includes a new AI Security dashboard that provides an overview of the AI models that are deployed in the cloud environment, what data they contain, and whether they are at risk. Orca covers risks end-to-end, from training and fine-tuning to production deployment and inference. Orca connects with your AI data and processes on all levels – the managed cloud AI services, the unmanaged AI models and packages that your developers are using, and even with shift-left detection of leaky secret tokens AI services in your codebase.

#1. AI and ML BOM + inventory

Challenge: Much like other resources in the cloud, shadow AI and LLMs are a major concern. In their excitement to explore all the opportunities that Generative AI brings, developers are not always waiting for IT approval before integrating Gen AI services into their flows. And security is probably not the first thing on their mind. This leaves security teams in the dark about which AI projects have been deployed in their environment, whether they contain sensitive data, and whether they are secure.

Orca solution: Orca scans your entire cloud environment and detects all deployed AI models, providing a full inventory and Bill of Materials (BOM). Orca detects projects on Azure OpenAI, Amazon Bedrock, Google Vertex AI, AWS Sagemaker and those using one of 50+ most commonly used AI software packages, including Pytorch, TensorFlow, OpenAI, Hugging Face, scikit-learn, and many more.

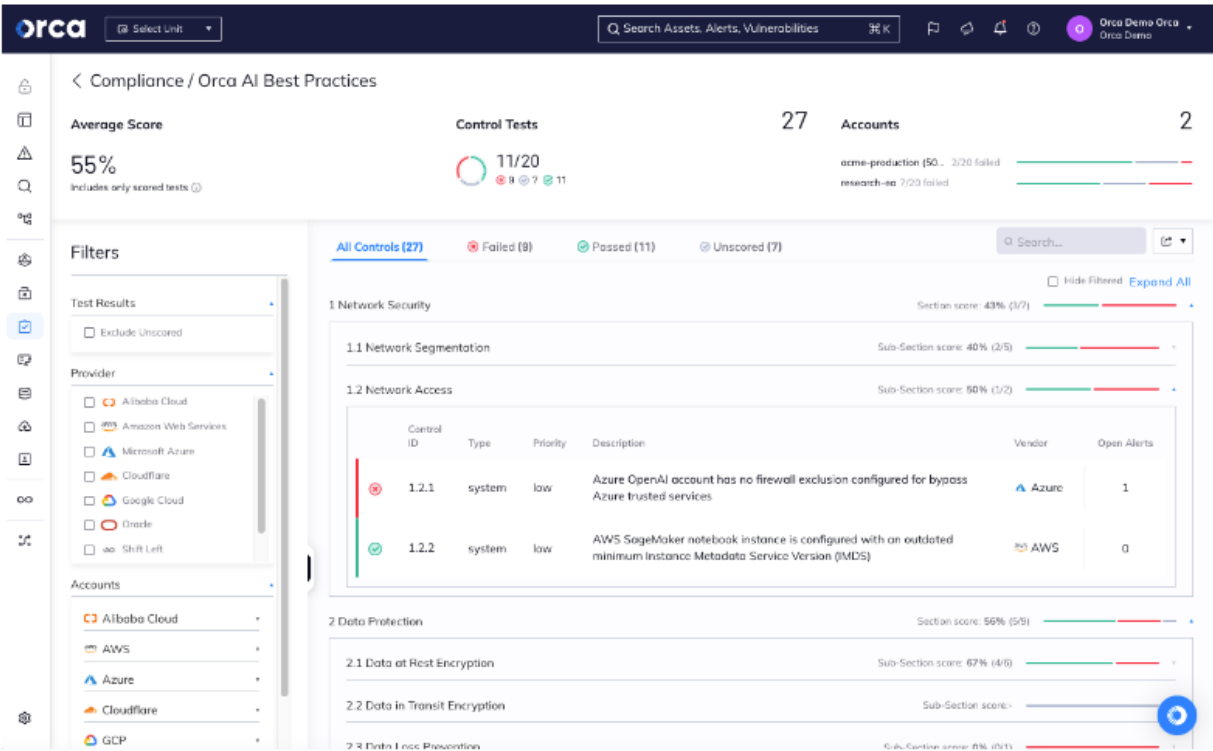


The Orca Platform shows a full inventory and bill of materials of all deployed AI models

#2. Warn when AI projects are publicly accessible

Challenge: Data used to train AI models is often sensitive and can contain proprietary information. Since it could be very damaging if this information got breached, it's very important to ensure that AI models are not publicly exposed. However, misconfigurations do happen, and especially with Shadow AI, security may not be the first thing developers have on their mind. This leaves security teams struggling to ensure that all AI models are being kept private.

Orca solution: Since Orca has insight into the AI model settings and network access, Orca will alert whenever public access is allowed, so security teams can quickly fix the issue to prevent any data breaches.

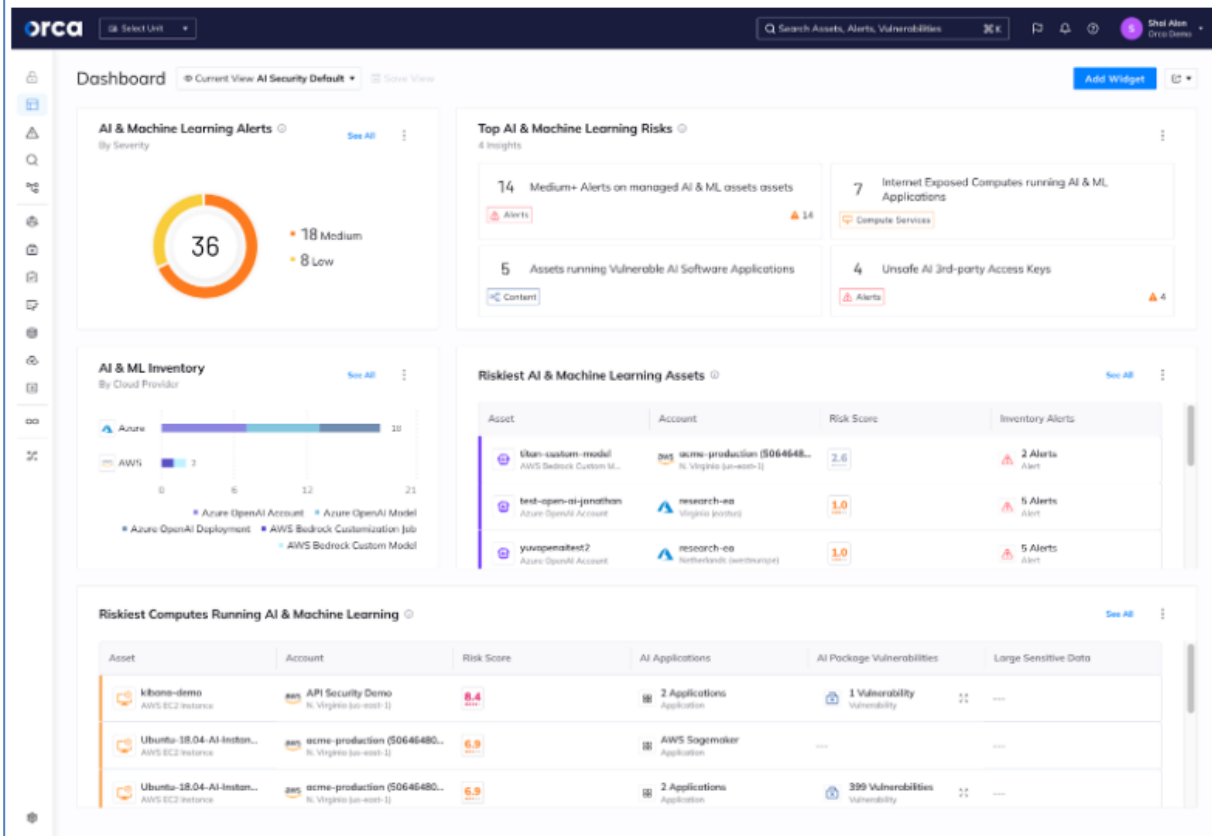


Orca AI Security best practices compliance performs AI-Security Posture Management

#3. Detect sensitive data in AI projects

Challenge: AI models use vast amounts of data to train models. It's possible that this data inadvertently contains sensitive data, and that developers and security teams are not even aware of the presence of the data, such as PII and access keys. LLMs and Generative AI models could then “spill out” this sensitive data, which could be used by bad actors. In addition, if developers are not aware that the training set contains sensitive data, they may not prioritize security of the resource as much as it needs to be.

Orca solution: Using Orca's [Data Security Posture Management \(DSPM\)](#) capabilities, Orca scans and classifies all the data stored in AI projects, and alerts if they contain sensitive data, such as email addresses, telephone numbers, email addresses, and social security numbers (PII), or personal health information (PHI). By informing security teams where sensitive data is located, they can make sure that these assets are protected with the highest level of security.



Orca displays pertinent security information in the AI Security dashboard

125. Upon information and belief, Orca further infringes multiple additional claims of the '529 patent. For example, but not limited to, claims 2, 4, 5, 6, 7 and 8.

126. Orca is aware of the '529 patent and Orca's infringement thereof at least as of the filing of these counterclaims. Moreover, this is another in a long line of examples of Orca copying Wiz, as discussed above. Accordingly, Orca has and continues to willfully infringe the '529 patent.

127. Orca has induced and continues to induce infringement of one or more claims of the '529 patent by encouraging customers to use AI-SPM in a manner that directly infringes those claims. Despite its knowledge of the existence of the '529 patent, since at least the filing of this Counterclaim, Orca, upon information and belief, continues to encourage, instruct, enable and otherwise cause its customers to use AI-SPM in a manner that infringes one or more claims of the '529 patent. Upon information and belief, Orca specifically intends that its customers use AI-SPM in a manner that infringes one or more claims of the '529 patent by, at a minimum, providing instructions and/or support documentation directing customers on how to use the AI-SPM in an infringing manner, in violation of 35 U.S.C. § 271(b). For example, Orca's public blog posts cited above provide instructions and encourage customers to practice all steps of the claimed method.

128. Orca has contributed and continues to contribute to the infringement of one or more claims of the '529 patent. Upon information and belief, Orca knows that AI-SPM features are especially made and/or adapted for users to infringe one or more claims of the '529 patent and is not a staple article or commodity of commerce suitable for substantial non-infringing use. Because Orca included features, such as for example but not limited to, AI-SPM in Orca's products, Orca intends for customers to use it. Upon information and belief, AI-SPM has no suitable use that is non-infringing, and therefore Orca intends for customers to use AI-SPM in an

infringing manner. Orca's sales of products including AI-SPM constitute contributory infringement in violation of 35 U.S.C. § 271(c).

PRAYER FOR RELIEF

WHEREFORE, Wiz respectfully asks that the Court enter judgment against Orca and in favor of Wiz as follows:

129. A judgment that Orca has infringed and continues to infringe (either literally or under the doctrine of equivalents) one or more claims of the Asserted Patents under at least 35 U.S.C. § 271(a);

130. A judgment that Orca has induced and continues to induce others to infringe one or more claims of the Asserted Patents under at least 35 U.S.C. § 271(b);

131. A judgment that Orca has contributorily infringed and continues to contribute to the infringement of one or more claims of the Asserted Patents under at least 35 U.S.C. § 271(c);

132. A judgment that Orca's infringement of the Asserted Patents has been and continues to be willful;

133. An award of monetary damages sufficient to compensate Wiz for Orca's patent infringement, with interest, pursuant to at least 35 U.S.C. § 284;

134. A preliminary and permanent injunction prohibiting Orca and its officers, agents, representatives, assigns, licenses, distributors, servants, employees, related entities, attorneys, and all those acting in concert, privity, or participation with them, from:

- (a) infringing or inducing the infringement of any claim of the Asserted Patents; and
- (b) soliciting any new business or new customers using any information or materials that Orca derived from its infringement of the Asserted Patents;

135. An award of enhanced damages of three times the amount found or assessed for Orca's willful patent infringement, pursuant to at least 35 U.S.C. § 284, including interest on such damages;

136. An order finding this case exceptional and awarding Wiz its attorneys' fees, to be obtained from any and all of Orca's assets, pursuant to 35 U.S.C. § 285, including prejudgment interest on such fees;

137. An accounting and supplemental damages for all damages occurring after the period for which discovery is taken, and after discovery closes, through the Court's decision regarding the imposition of a permanent injunction;

138. An award of Wiz's costs and expenses of this suit as the prevailing party; and

139. Any and all other relief in Wiz's favor that the Court deems just and proper.

JURY DEMAND

Orca hereby demands a trial by jury on all issues so triable.

OF COUNSEL:

Jordan R. Jaffe
Catherine Lacy
Callie Davidson
Alex Miller
WILSON SONSINI GOODRICH & ROSATI, P.C.
One Market Plaza
Spear Tower, Suite 3300
San Francisco, CA 94105
(415) 947-2000

Praatika Prasad
Wilson Sonsini
1301 Avenue of the Americas, 40th Floor
New York, NY 10019-6022
(212) 999-5800

Dated: June 4, 2024

/s/ Kelly E. Farnan

Frederick L. Cottrell, III (#2555)
Kelly E. Farnan (#4395)
Christine D. Haynes (#4697)
RICHARDS, LAYTON & FINGER, P.A.
One Rodney Square
920 N. King Street
Wilmington, DE 19801
(302) 658-6541
cottrell@rlf.com
farnan@rlf.com
haynes@rlf.com

Counsel for Defendant Wiz, Inc.